ESD-TR-89-308 MTR-10568

Network Repair Level Analysis Program: User's Manual

By

R. A. Moynihan
M. A. Strauch
A. E. Taub

December 1989

Prepared for
Director of Logistics, Joint STARS Joint Program Office
Electronic Systems Division
Air Force Systems Command
United States Air Force
Hanscom Air Force Base, Massachusetts

Project No. 646C
Prepared by
The MITRE Corporation
Bedford, Massachusetts
Contract No. F19628-86-C-0001

REVIEW AND APPROVAL

This technical report has been reviewed and is approved for publication.

PETE PETERSON
Chief, Supply Support Division
Directorate of Logistics
Joint STARS Joint Program Office

FOR THE PROGRAM MANAGER

STEPHEN H. FARISH, LT COL, USAF
Director of Logistics
Joint STARS Joint Program Office

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified	1b. RESTRICTIVE MARKINGS
2a. SECURITY CLASSIFICATION AUTHORITY	3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited.
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE	
4. PERFORMING ORGANIZATION REPORT NUMBER(S) MTR-10568 ESD-TR-89-308	5. MONITORING ORGANIZATION REPORT NUMBER(S)
6a. NAME OF PERFORMING ORGANIZATION The MITRE Corporation 6b. OFFICE SYMBOL *(If applicable)*	7a. NAME OF MONITORING ORGANIZATION
6c. ADDRESS *(City, State, and ZIP Code)* Burlington Road Bedford, MA 01730	7b. ADDRESS *(City, State, and ZIP Code)*
8a. NAME OF FUNDING/SPONSORING ORGANIZATION Director of (continued) 8b. OFFICE SYMBOL *(If applicable)* ESD/JSL	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER F19628-86-C-0001
8c. ADDRESS *(City, State, and ZIP Code)* Electronic Systems Division, AFSC Hanscom AFB, MA 01731-5000	10. SOURCE OF FUNDING NUMBERS PROGRAM ELEMENT NO.: — PROJECT NO.: 646C TASK NO.: — WORK UNIT ACCESSION NO.: —

11. TITLE *(Include Security Classification)*

Network Repair Level Analysis Program: User's Manual

12. PERSONAL AUTHOR(S)

Moynihan, R. A., Strauch, M. A., Taub, A. E.

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT *(Year, Month, Day)*	15. PAGE COUNT
Final	FROM ____ TO ____	1989 December	76

16. SUPPLEMENTARY NOTATION

17. COSATI CODES			18. SUBJECT TERMS *(Continue on reverse if necessary and identify by block number)*
FIELD	GROUP	SUB-GROUP	Maintenance Cost Trade-Off Model Network Flow Algorithm Network RLA (NRLA) (continued)

19. ABSTRACT *(Continue on reverse if necessary and identify by block number)*

Every major Air Force acquisition program has a requirement to make the most cost-effective repair level decisions for all system components. The MITRE Network Repair Level Analysis (NRLA) program calculates these best decisions on the basis of input cost factors provided by a separate program-specific life cycle cost (LCC) model, such as the Joint STARS LCC model. In contrast, the standard AFLC NRLA Model incorporates a built-in, generic LCC model which makes simplifying assumptions that do not apply to many Air Force systems, e.g., that equipment deployments and maintenance capabilities at all base locations are identical. Thus the MITRE NRLA program provides the capability to determine repair level decisions based on more accurate input cost factors in cases where a program-specific LCC model exists.

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT ☐ UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT. ☐ DTIC USERS	21. ABSTRACT SECURITY CLASSIFICATION Unclassified
22a. NAME OF RESPONSIBLE INDIVIDUAL Judith Schultz	22b. TELEPHONE *(Include Area Code)* (617) 271-8087 22c. OFFICE SYMBOL Mail Stop D135

DD FORM 1473, 84 MAR

83 APR edition may be used until exhausted.
All other editions are obsolete.

SECURITY CLASSIFICATION OF THIS PAGE

UNCLASSIFIED

8a. Logistics, Joint STARS Joint Program Office

18. Optimum RLA (ORLA)
Repair Level Analysis (RLA)

ACKNOWLEDGMENTS

This document has been prepared by The MITRE Corporation under Project No. 646C, Contract No. F19628-86-C-0001. The contract is sponsored by the Electronic Systems Division, Air Force Systems Command, United States Air Force, Hanscom Air Force Base, Massachusetts 01731-5000.

The authors appreciate the support of the Joint STARS Program Office in developing the NRLA Program, which should serve as a useful tool on many other Air Force programs as well.

They also wish to thank Lt. Hank Birkdale of ESD/PLLC for several valuable discussions regarding methods to implement an independent NRLA algorithm. In addition, B. K. Schmidt provided analytical support in investigating alternative solution algorithms. Word processing and other manuscript preparation tasks were efficiently performed by J. B. MacGillivray.

TABLE OF CONTENTS

SECTION	PAGE

LIST OF ILLUSTRATIONS

LIST OF TABLES

SECTION 1

PURPOSE OF MITRE NRLA PROGRAM

The MITRE Network Repair Level Analysis (NRLA) program is intended to serve as a supplementary tool to the current NRLA model [1] which was implemented and is maintained by Air Force Logistics Command (AFLC). Both of these software packages are intended to calculate most cost-effective repair level decisions (i.e., repair at base, repair at depot, or scrap on failure) for all line replaceable units (LRUs) and shop replaceable units (SRUs) within a given Air Force system. To make such decisions, both models must obtain specific LRU, SRU and support equipment (SE) costs associated with all possible repair options.

In order to provide a self-contained repair level analysis (RLA) capability, the AFLC NRLA incorporates a generic LCC model which will calculate the particular costs which must be provided as input to the RLA algorithm. Thus, the AFLC Model is actually a combined life cycle cost (LCC) model and RLA algorithm. This feature is very convenient, especially in cases where no prior LCC model or alternative cost calculation method exists.

However, since the LCC model component of the AFLC NRLA is generic, a number of simplifying assumptions were made in its structure in order to facilitate computer implementation. For example, the equations of this LCC model treat all deployment bases as being identical, both in terms of maintenance capabilities and in quantities and type of deployed equipment.

For many Air Force programs, e.g., the Joint STARS and MILSTAR Terminal Segment Systems, this assumption simply does not hold. In fact, there is often wide variation in equipment deployment from base to base with some bases providing greater maintenance capabilities than others. For example, under both the Joint STARS and MILSTAR support concepts some bases will be designated as centralized intermediate maintenance facilities (IMFs) which will provide repair support to the remaining base locations. Therefore, even if all LRUs and SRUs are designated as base repairable, only the IMFs will require the associated support equipment. In addition, the calculation of initial spares requirements will vary significantly from base to base, especially in cases where bases receive not only different quantities, but different <u>types</u> of equipment. For example, the MILSTAR system might include 100 B-52 terminals which are deployed to only 4 bases out of a total of 70 MILSTAR bases. Therefore, only these 4 bases will require spares of B-52 terminal-specific LRUs and SRUs. These effects do not average out when using a uniform base deployment assumption. Thus, application of the AFLC NRLA model to systems which do not satisfy this uniformity of deployment bases assumption will necessarily lead to inaccuracies in the LCC calculations, which in turn will create uncertainty as to the cost-effectiveness of the resulting repair level decisions.

In some cases, again such as in the Joint STARS and MILSTAR programs, a detailed, system-specific LCC model [2,3] does exist which represents the system deployment and logistics structure much more accurately. However, the computer implementation of the AFLC NRLA model does not permit its two components, the generic LCC model and the RLA algorithm, to be separated and executed independently. In particular, its generic LCC model cannot be bypassed so that LCC inputs can be provided directly to the RLA algorithm.

The MITRE NRLA program provides exactly this last capability. It executes the same network RLA algorithm as the AFLC model does, but it allows for -- in fact, it requires -- direct input of the required LCC factors. These LCC factors may be generated by any chosen LCC model or other calculation method, and need only be provided in properly formatted data files. Thus the MITRE NRLA program may be applied when the assumptions of the AFLC generic LCC model do not hold and there is an alternative means of generating the required LCC inputs to the NRLA algorithm.

The AFLC NRLA user's guide provides a thorough discussion of the RLA problem and the network flow algorithm used to solve it. Since we wish to keep our own user's manual reasonably self-contained, we will also provide a discussion on these topics. This redundancy is for the reader's convenience. In fact, we have made use of a number of well-phrased paragraphs from the AFLC document. These borrowed paragraphs include sections 2.1, 2.2, and 2.5. In addition, appendix A is a slightly edited version of section 4 in the AFLC guide.

SECTION 2

NETWORK REPAIR LEVEL ANALYSIS ALGORITHM

2.1 OVERVIEW OF REPAIR LEVEL PROCESS

Before an Air Force system is deployed in the field, repair-level decisions must be made for the component LRUs and SRUs of the end item being procured. As a prerequisite to these decisions, each component must be analyzed to determine the different ways in which it may be unable to perform its required functions. Inability to perform a required function constitutes a failure. A component failure may occur when an LRU fails as the result of the failure of one or more of its internal SRUs, or, as a result of a failure which is not within a specific SRU item. The LRU/SRU or LRU/non-SRU failure combination is called a "failure mode". For those failure modes where repair is technically feasible, a repair level decision must be made where the decision involves choosing among depot level repair, discard (scrap), and intermediate (base) repair. For many LRUs, all repairable failure modes are assigned the same repair level alternative; however, some LRUs could have intermediate repair assigned for some failure modes and depot repair assigned for other failure modes. Included among the factors influencing each repair level decision are the life cycle costs associated with each repair level and the availability of support equipment and repair facilities at base and depot.

The NRLA solution method described below deals with economic factors affecting the repair level decisions. In the model, the economic analysis of repair level decisions is based on specific life cycle costs associated with each repair level option. The NRLA calculation need not address a comprehensive set of LCC elements. It must include only those elements which directly impact the repair level decision.

2.2 THE RLA NETWORK

The NRLA solution algorithm is an analytic optimization technique which may be applied to making repair level decisions. This technique is fundamentally different from the previous Air Force repair level analysis methodology in a number of significant ways. The previous Air Force RLA methodology [4] was a heuristic procedure which tried to account for the costs of shared resources, such as support equipment, by distributing their costs, using a proration scheme, to the individual LRUs and SRUs that required them. This proration procedure did not always produce the least total cost solution for the overall system. In fact, for certain examples, it was shown to recommend the <u>highest</u> cost set of decisions. The network algorithm approach to the RLA problem was developed by the MITRE Corporation[1] and was subsequently adopted by AFLC as the standard RLA methodology [1].

The NRLA algorithm is a true optimization technique which eliminates the key faults of the previous heuristic RLA procedure. It determines repair level recommendations simultaneously

[1]The first author should be contacted for further information on this early work.

for all the failure modes of a group of LRUs and for the SRUs associated with the LRUs, as opposed to considering each failure mode and SRU independently. The support equipment required to accomplish LRU and/or SRU repair is considered to be a resource whose cost must be economically justified by the group of LRUs and/or SRUs which require it. This cost is jointly shared by the group of items requiring the resource; however, the cost is not prorated to the individual items in the group. Finally, as a consequence of these fundamental differences, the model is able to make repair level recommendations such that the total cost for the group of LRUs and SRUs is minimized. Thus, this method determines the economically optimal set of repair level decisions for the entire group of items.

As previously stated, the solution algorithm determines repair level recommendations based solely on economic considerations. Further, the costs associated with each repair level option should be those which are specifically incurred as a result of choosing the option. Thus, the life cycle costs for spares and support equipment are included because the total expenditure is a function of the repair level decisions.

Cost values used by the model for determining repair level decisions should be based on data factors relevant to the

(1) End-item utilization
(2) Maintenance system costs
(3) Supply system costs
(4) Support equipment costs, and
(5) LRU and SRU repair costs.

Both the MITRE and AFLC versions of the NRLA model formulate the repair level decision problem in terms of a network. This approach is used because it specifically considers LRU to SRU indenture level relationships and it treats each unique piece of SE as a repair resource which is shared by a group of LRUs and/or SRUs. This model uses life cycle costs with LRU/SRU/SE interdependency relationships to construct a network representation of the repair level decision problem. The network contains nodes for the LRU failure modes, the associated SRU's and the SE resources required for LRU and/or SRU repair. As illustrated in figure 2-1 below, the nodes are connected by arcs representing the costs for the different repair level options. A network analysis optimization algorithm, known as *Max-Flow Min-Cut*, is applied to this network to determine the minimum cost set of repair level decisions. (See reference 5 for an explanation of this algorithm.)

The DEC1 and DEC2 arcs shown in figure 2-1 represent SRU costs that are incurred with different combinations of LRU and SRU repair level decisions. These costs are discussed in the next section.

It should be emphasized that figure 2-1 represents the network component for a *single* LRU failure mode which is associated with a particular SRU failure. The full network will contain the components associated with all LRU failure modes, each connected to the same source (S) node and sink (T) node. A repair level decision must be made for each LRU/SRU failure mode combination. Thus cost inputs are required for all arcs of this network for all LRU failure modes.

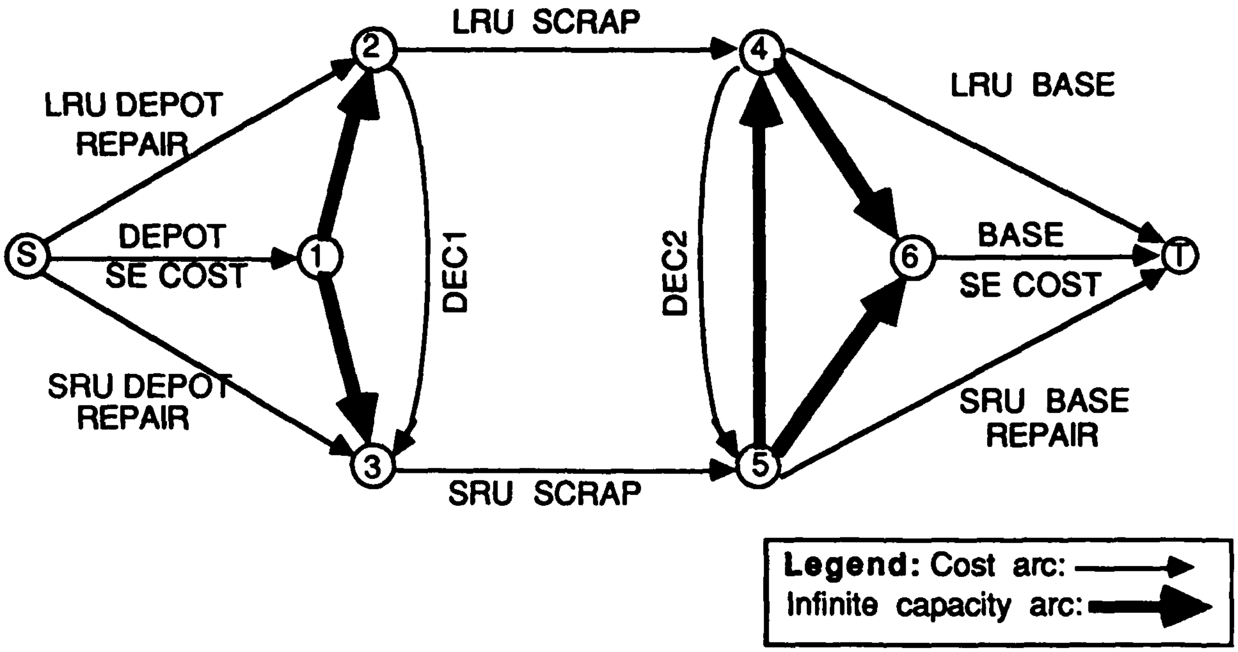

Figure 2-1. Basic NRLA Network

An elementary discussion on how maximum flow in a network corresponds to a minimum cost repair level decision is given in section 2.5. More details on formulating an RLA network and solution methods are given in appendices A and B.

2.3 COMPUTATION OF COST COMPONENTS FOR NETWORK ARCS

The AFLC NRLA model uses its built-in LCC equations to calculate all the required cost inputs for the network arcs. These LCC calculations use basic LRU and SRU input parameters, such as mean time between failure (MTBF), unit cost, repair man-hours per failure, and labor rates. However, as discussed in section 1, the MITRE NRLA program requires that these network arc costs be provided as direct inputs, i.e., they must be computed outside the NRLA program. This section provides some guidelines to aid in the computation of these cost inputs.

From our discussion in section 2.2, these costs must be computed for each failure mode of a given LRU. For example, an LRU could have six different failure nodes, five associated with internal SRUs and one not. This LRU would then generate six of the network components shown in figure 2-1 above, with the last one containing no SRU, DEC1, or DEC2 arcs.

Thus, to fill out the costs on the arcs of these network components, the LRU costs for each repair decision (depot repair, scrap, and base repair) must be calculated and apportioned to the six different failure modes. This apportionment may be carried out, as done in the AFLC NRLA model, by dividing the LRU cost in proportion to the failure rate of the internal SRUs associated with the LRU failure modes.

In general, the sum of the failure rates of internal SRUs cannot logically be greater than that of the parent LRU. In addition, if the sum of the SRU failure rates is less than the LRU failure rate, then there must be an LRU failure mode which is not associated with an SRU. By such considerations one can then allocate LRU costs to LRU failure modes.

Continuing the example above, suppose that the total LRU depot repair cost is $20,000, that the failure rate of the LRU is 100 failures per million operating hours (fpm), and that the five internal SRUs have respective failure rates of 8, 10, 15, 20, and 30 fpm. We then note that the non-SRU-related failure mode of the LRU must have a failure rate of 17 fpm. The depot repair costs for the six LRU failure modes may then be calculated as $1,600, $2,000, $3,000, $4,000, and $6,000 for the five respective SRU-related failure modes, and $3,400 for the non-SRU-related failure mode. (Note that these figures are for repair of the LRU only. SRU repair costs will also have to be computed separately.)

In applying the above proration scheme in practice, one must be careful to correct for bad data. It is not uncommon to receive LCC data where the sum of the SRU failure rates is <u>greater</u> than the parent LRU failure rate. In such cases, one must either correct the errors or, at least, renormalize the SRU failure rates before allocating the LRU cost. If such a correction is not made then the sum of the costs allocated to the LRU failure modes will be greater than the total LRU cost.

Also note that, since an SRU can be associated with only one LRU failure mode, there is no similar cost allocation problem for SRU costs. However, a portion of the SRU costs must be placed on the DEC1 and DEC2 arcs of the network, as shown in table 2-1 and described in detail in section 2.4 below.

Table 2-1 shows the breakout of logistics costs which the AFLC NRLA model includes in the respective repair level decision network arcs. This table may also be used as a guideline in preparing cost inputs to the MITRE NRLA program. In general, all LRU or SRU costs which are incurred under a given repair level option should be included in the cost capacity of the associated network arc. Thus table 2-1 should be interpreted as identifying those logistics costs which are usually incurred under each option.

For example, under the LRU "Base Repair" option nearly all costs are incurred. However, the Depot Initial Spares cost is not included since the LRU is assumed to be always repaired at the base. Thus the depot never enters into the repair pipeline. Also, under the "Base Repair" option, no Replacement Spares are required because the LRU is never condemned after failure. A detailed LCC model, such as the Joint STARS model [2], would automatically zero out most cost elements that do not apply to a particular repair level decision, but table 2-1 can serve as a checklist to ensure that this has happened (e.g., that Training costs are not included under "Scrap" costs).

Table 2-1. RLA Network Cost Components

	LRU Costs			SRU Costs					SE Costs	
Logistics Costs	Depot Repair	Scrap	Base Repair	Depot Repair	Scrap	Base Repair	DEC1	DEC2	Depot Repair	Base Repair
SE Acquisition & Maintenance									X	X
Technical Data Acquisition	X		X	X		X				
Maintenance Training	X		X	X		X				
Repair Labor	X		X	X		X				
Inventory Item Management	X	X	X	X	X	X				
Repair Materials	X		X	X		X				
Packing & Shipping	X	X					X	X		
Base Initial Spares	X	X	X			X		X		
Depot Initial Spares	X			X			X			
Replacement Spares		X			X					
Total	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10

Two exceptions to including all incurred LRU or SRU costs on the associated network arcs are (1) the support equipment costs which belong on their own separate arcs, and (2) the DEC1 and DEC2 costs which are SRU costs resulting from specific combinations of LRU/SRU repair level decisions, as shown in table 2-2.

Note that if a "Scrap" decision is made for an LRU in an SRU-related failure mode, then the SRU costs "disappear" in the sense that when an LRU is discarded the (failed) SRU is discarded with it. Thus the LRU "Scrap" costs cover the Replacement Spares cost of scrapping the SRU, and clearly there is no need of SRU Initial Spares or separate SRU packing and shipping.

However, SRU Replacement Spares costs are incurred if the LRU is repaired and the SRU is scrapped, and the SRU Initial Spares and Packing and Shipping costs in this case will change depending on whether the LRU is repaired at the base or at the depot. In fact, even when the SRU is depot repaired, these latter two costs elements will change as a function of the LRU repair location. To handle all these situations the SRU costs must be allocated to the three arcs: SRU Scrap, DEC1 and DEC2.

First, table 2-1 indicates that only the SRU Replacement Spares cost should be included on the "SRU Scrap" arc, shown in figure 2-1. Note that table 2-2 indicates that, because of the way the network algorithm works, the cost of the LRU "Scrap" decision (number 3) is determined by adding the costs of both the LRU and SRU Scrap arcs. The MITRE NRLA program handles this problem internally by first subtracting this SRU Replacement Spares cost from the LRU Scrap cost and placing the difference on the LRU Scrap arc. (Thus when the algorithm adds these two arcs it gets only the LRU scrap cost without double-counting by including SRU Replacement Spares.). The NRLA program user should therefore input the full LRU Scrap cost, including the cost elements identified in table 2-1, without worrying about this subtraction process. We point out that, in this situation, bad data can lead to a negative arc cost if the SRU Scrap cost is larger than the LRU Scrap cost. Of course, logically, such a circumstance should never occur.

Next, note from table 2-2 that under the "LRU Base Repair, SRU Scrap" option (number 5) that DEC2 costs are now added to the SRU Scrap costs. In this situation, since the SRU Scrap arc accounts for the SRU Replacement Spares cost only, the DEC2 arc must cover the cost of SRU Base Initial Spares (SRU stock is now required for base LRU repair) and the one-way Packing and Shipping of Replacement Spares from depot to base.

Finally, under the "LRU Base Repair, SRU Depot Repair" option (number 6), table 2-2 indicates that SRU costs include the SRU Depot Repair, DEC1 and DEC2 arcs. As we will discuss further below, the SRU Depot Repair cost should be computed under the assumption that its higher LRU is also Depot Repaired. Thus the SRU Depot Repair arc cost does not include any Base Initial Spares or Packing & Shipping costs. In addition, the SRU Depot Initial Spares in this case (LRU Depot Repair) would be computed using a shorter depot repair cycle time, since the failed SRU is not removed until the LRU is repaired at the depot. Therefore, when the LRU is Base Repaired and the SRU is Depot Repaired we need to pick up these additional costs via the DEC1 and DEC2 costs. As mentioned, the DEC2 cost includes Base Initial Spares and one-way Packing & Shipping. Thus DEC1 must include a second one-

Table 2-2. RLA Decision Costs

RLA Decision Costs						
	Decision		Included Costs			
			Support Equipment			
	LRU	SRU	Depot	Base	DEC1	DEC2
1	Depot	Depot	Yes	No	No	No
2	Base	Base	No	Yes	No	No
3	Scrap	Scrap	No	No	No	No
4	Depot	Scrap	Yes	No	No	No
5	Base	Scrap	No	Yes	No	Yes
6	Base	Depot	Yes	Yes	Yes	Yes
7*	Scrap	Depot	Yes	No	Yes	No

*This anomalous alternative cannot be eliminated from the network for technical reasons. However, it is extremely unlikely to occur. Methods of handling it should it occur are given in appendix A.

way Packing & Shipping cost (to cover transportation of the failed SRU from base to depot) and the increase in Depot Initial Spares resulting from a longer depot repair cycle time. (The depot now has to wait for the failed SRU to be shipped from the base before it enters the depot maintenance queue.)

The NRLA program applies the user-specified LRU/SRU SE requirements (given in input file 5 as discussed in section 4.2.3 below) to determine which SE costs are incurred by a particular repair level decision. This process is straightforward; if an LRU or SRU is depot repaired then the NRLA algorithm will include the cost of any depot SE required. The same is true for base repair. However, as mentioned in section 2.2, the cost of any required SE will be shared among all LRUs/SRUs that use it. Thus, table 2-2 should be interpreted to imply that Base and/or Depot SE costs are incurred by whatever LRU/SRU repair decisions that require them.

Further insight into how these costs should be calculated is provided by the LCC equations given in section 3 of the AFLC NRLA Model User's Guide. The 10 costs (T1 through T10) given in table 2-1 are associated with the respective LRUs, SRUs, and SEs. The NRLA algorithm then selects, for each LRU failure mode and associated SRU, one of the six decisions illustrated in table 2-2. (Unfortunately, the seventh decision, an anomalous situation, can also be occasionally selected and must be forced out, as discussed in appendix section A.6). The algorithm selects a total set of repair decisions that will yield the least total cost for the entire system. The same network is used for sensitivity analysis investigations by properly modifying the network arc costs.

It is possible, as we have done for the Joint STARS program, to modify an existing LCC model so that it will automatically calculate these required network arc costs. The Joint STARS LCC model will also output these costs to a properly formatted computer data file, ready for use by the MITRE NRLA program.

2.4 USING AN LCC MODEL TO COMPUTE NETWORK COSTS

Several different computer runs of any LCC model will be required to generate the required inputs to the NRLA program, because these calculations must be made based on several different LRU and SRU repair level choices. Below we describe a sequence of four computer runs of an appropriate LCC model (or other calculation means) which will efficiently generate the required cost on the network arcs shown in figure 2-1.

LCC Model Run #1:

This computer run should be made with the assumption that all LRUs and SRUs are repaired at the depot. For each LRU failure mode, all LRU-specific costs should be accumulated in order to calculate the total cost "T1" shown in table 2-1. This calculation will probably require an apportionment of total LRU costs to the various failure modes in proportion to the relative failure rates of the associated SRUs, as discussed in section 2.3. Similarly, for each SRU, all SRU-specific costs should be summed to determine the cost total "T4" shown in table 2-1. For each LRU failure mode and associated SRU, these paired T1 and T4 costs are associated, respectively, with the "LRU Depot Repair" and "SRU Depot Repair" network arcs shown in figure 2-1. (See the discussion below concerning the allocation of LRU costs to failure modes.)

These paired T1 and T4 costs should be listed on successive lines in NRLA Input Data File 1, in accordance with the format described in section 4.2.1. For sensitivity analysis purposes the T1 and T4 costs, as well as all other network arc costs, must also be broken down into five specific components, as described in section 3.

LCC model Run #1 should also be used to compute the acquisition and maintenance costs for depot SE. For each different depot SE type, this cost should be computed to represent the total "T9" shown in table 2-1. These T9 costs are associated with the "Depot SE Cost" arcs in figure 2-1 and should be entered in NRLA Input Data File 6b in the format described in section 4.2.5.

To compute depot SE costs, the LCC model must utilize data specifying the depot SE requirements for repair of each LRU and SRU. Similar LCC model data must exist to specify SE requirements for base level repair in LCC model run #2, described below. Together, this LRU/SRU/SE cross reference data should be used to fill out NRLA Input Data File 5, as described in section 4.2.3. Whether File 5 is created manually or printed out by the LCC model is the user's decision.

Caution: Note that for Run #1 to make sense, the LCC model input data must include all LRU and SRU requirements for depot repair. Thus all SE, training, Tech Data, repair times, and repair materials factors must be specified in the LCC model data base. A similar comment applies to Run #2 with regard to base level repair.

LCC Model Run #2:

This computer run should be made with the assumption that all LRUs and SRUs are repaired at the base level. For each LRU failure mode, all LRU-specific costs should be accumulated in order to calculate the total cost "T3" shown in table 2-1. Similarly, for each SRU, all SRU-specific costs should be added to determine the cost total "T6" shown in table 2-1. For each LRU failure mode and associated SRU, these paired T3 and T6 costs are associated, respectively, with the "LRU Base Repair" and "SRU Base Repair" network arcs shown in figure 2-1. These paired T3 and T6 costs should be listed on successive lines in NRLA Input Data File 2, in the format specified in section 4.2.1.

LCC model Run #2 should also be used to compute the acquisition and maintenance costs for base level SE. For each different base level SE type this cost should be computed to represent the total T10 shown in table 2-1. These T10 costs are associated with the "Base SE Cost" network arcs in figure 2-1 and should be entered in NRLA Input Data File 6a in the format described in section 4.2.4.

LCC Model Run #3:

This LCC computer run should be made under the assumption of LRU scrap, i.e., discard on failure. Note that, under this assumption of LRU scrap, the SRU costs should disappear since SRUs are never removed from the failed LRUs. Depending on how the user's LCC model works, however, it may be wise to designate the SRUs as being discard-on-failure. It may also be good policy on this particular LCC calculation to zero out any LCC model input parameters that represent costs which will not be incurred under an LRU scrap decision. For example, maintenance training and technical data requirements should probably be zeroed out so that they are not inadvertently included in the calculated LRU scrap cost.

Even though no SRU costs should be generated by Run #3, the LRU-specific costs must still be kept track of in terms of LRU failure modes, many of which may be associated with specific SRUs. Thus, for each LRU failure mode, all LRU-specific costs should be added to determine the cost total "T2" shown in table 2-1. These T2 costs are associated with the "LRU Scrap" network arc shown in figure 2-1.

Despite the fact that no SRU-specific costs are gathered by Run #3, the T2 costs for each LRU failure mode must still be paired with any associated SRU when they are listed in NRLA Input Data File 3. See section 4.2.1 for the specific data file format requirements.

LCC Model Run #4

This LCC computer run is made to calculate the remaining network arc costs; that is, SRU Scrap, DEC1 and DEC2 costs. To make these calculations LRUs should be designated as repaired at the base level and SRUs should be specified as discard on failure (i.e., scrap). As mentioned above, all LCC model input data should be properly initialized, or zeroed out, to accurately represent these respective LRU/SRU repair options.

Only SRU costs are required from Run #4. For each SRU type, the cost of Replacement (or Condemnation) Spares generated by this run should be used to determine the "T5" cost shown in table 2-1. Each T5 cost is associated with the respective "SRU Scrap" arc shown in figure 2-1.

Also from this run, the one-way Packing & Shipping Cost (to send replacement spare SRUs from depot to base) and the cost of SRU Base Initial Spares (required since the LRUs are base repaired) should be calculated for each SRU and assigned as the "T8" cost shown in table 2-1. Each T8 cost then corresponds to the "DEC2" arc shown in figure 2-1.

Calculation of the final DEC1 cost may require a couple of special equations not normally executed in a LCC computer run with LRUs base repaired and SRUs scrapped (as in Run #4). As discussed previously, the DEC1 cost must include a second one-way Packing & Shipping cost and the increase in Depot Initial Spares resulting from a longer SRU depot repair pipeline. (See the discussion in section 2.3.)

The one-way SRU Packing & Shipping cost should be naturally calculated by LCC Run #4 and may easily be assigned to DEC1 in addition to DEC2. However, the increased SRU Depot Initial Spares is not a natural product of Run #4, or any of the other LCC computer runs we have described. In fact, for most standard LCC models, Run #4 would probably produce a zero cost for SRU Depot Initial Spares, since each SRU is a scrap item.

Thus, the increased Depot Initial Spares component of DEC1 is really a stand-alone special calculation. It should be calculated assuming that the SRU is a depot-repairable item. A standard depot sparing algorithm should be applied twice, once assuming that the failed SRU is removed (from its higher LRU) at the base level, so that the depot repair cycle time (DRCT) is approximately 50 days - the Air Force CONUS standard value from AFLCP 173-10, and a second time assuming that the failed SRU is removed at the depot (i.e., assuming its higher LRU is depot repaired), so that the DRCT value is smaller, perhaps only 20-25 days. The difference in these two calculations would then equal the increased Depot Initial Spares cost which should be added to the DEC1 cost.

The resulting total DEC1 cost then corresponds to the "T7" cost shown in table 2-1, and is assigned to the "DEC1" arc shown in figure 2-1. After the SRU Scrap, DEC1 and DEC2 costs have been computed for each SRU, they should be listed in NRLA Input Data File 4 in the format described in section 4.2.2.

2.5 MAX-FLOW MIN-CUT APPLICABILITY

The NRLA model employs a max-flow min-cut algorithm to determine repair level decisions which minimize expected costs. The applicability of this technique will be explained by an example.

Consider a situation in which a fluid is to be pumped through a pipeline from location S to location T. The locations are connected, via intermediate locations, by two pipelines, S-A-C-T and S-B-D-T as shown in figure 2-2. The numbers next to each arc represent the fluid carrying

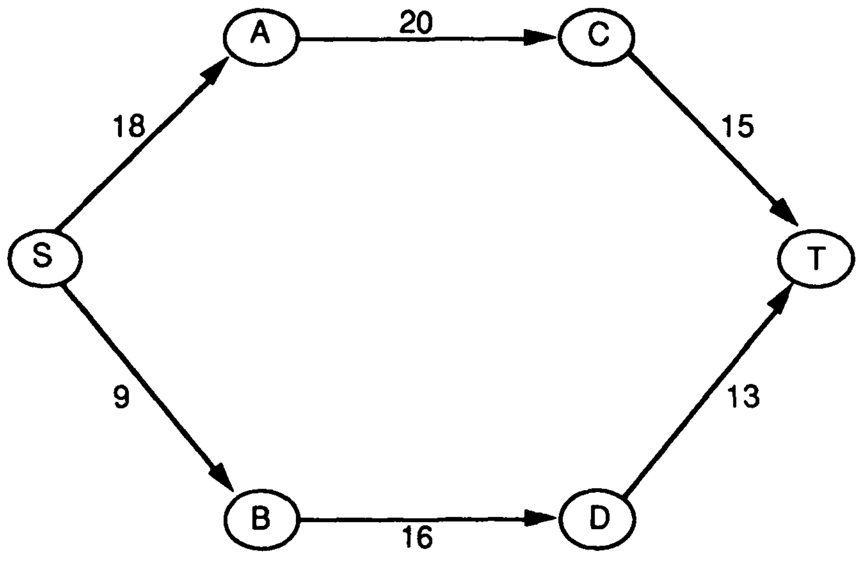

Figure 2-2. Sample Network.

capacity of the pipe connecting the locations. Thus, the maximum flow from location S to location A is 18 gallons per minute and from S to B the maximum is 9. Therefore, the total output capacity of S is 18 plus 9, or 27. Similarly, at T the maximum input capacity is 15 plus 13, or 28. However, the actual maximum flow from S to T is not 27 or 28. The flow along the route S-A-C-T is constrained to 15, the maximum capacity from C to T. Similarly, the maximum capacity from S to B limits the flow along the other pipeline to 9. Consequently, the total maximum flow from S to T is 15 plus 9, or 24.

The above discussion focused on finding the maximum flow through a network. The applicability of this maximization process may not seem obvious for the repair level decision problem in which the objective is to find a minimum cost solution. However, let us take a different view of the network shown in figure 2-2.

Suppose that the network still represents a pipeline through which a fluid flows but the problem to be solved is to determine the best way to stop the fluid flow. Assume that the flow in any of the sections of the pipeline can be stopped by purchasing a plug, and that the cost of a plug is directly proportional to the capacity of the pipe for which it is purchased. For example, a plug for the section from S to A would cost $18, and for the section from S to B, $9.

The total flow could be halted by plugging the two pipes emanating from S at a cost of $18 plus $9, or $27. Similarly, the flow could be halted at T by purchasing plugs costing $15 and $13, or $28. There are seven other combinations of two plugs which can be used to cut off the flow from S to T. The least cost combination involves the arcs S-B and C-T at a total cost of $24.

The fact that the numeric value of the minimum cut, 24, is identical to the value for the maximum flow is not coincidental. This occurs because each is directly a function of the constraining arcs in the network. That is, by finding the minimum cut for a network the maximum flow is also found.

To convert from the fluid flow network to an RLA cost network, consider a redefinition for the numbers on the arcs. Now the 18 and 9 represent the cost of performing depot repair for two different items. The 20 and 16 represent the costs for scrapping the items, while 15 and 13 represent base repair costs. The problem is now to find the least-cost repair alternatives.

The cost for depot repair of both items would be 18 plus 9, or 27; scrapping both items would be 36; and base repair for both items would be 28. The least-cost alternative is 15 plus 9, or 24. This decision represents base repair for one item and depot repair for the other. Note that the cost of these two decisions, 24, is the same as the result obtained in the fluid flow example when considering the maximum flow and the minimum cut. Thus, the repair level decision problem can be formulated in terms of a network and solved with a max-flow min-cut algorithm to determine the last cost decisions.

Although the above example does not include LRU to SRU relationships or support equipment requirements they are easily accommodated into the network structure. The method for including them is presented in appendix A.

SECTION 3

SENSITIVITY ANALYSIS

3.1 PURPOSE OF SENSITIVITY ANALYSIS

The MITRE NRLA program may be utilized to calculate the impact on LRU/SRU repair level decisions that result from changes in the original item unit cost and item MTBF. This sensitivity analysis is typically performed to see how stable the RLA decisions are. Since there is always uncertainty in unit cost and MTBF estimates made during system development, it is important to know how large an error or change in these estimates would have to be in order to change particular RLA decisions. In addition, other program considerations might override RLA decisions which are borderline choices with regard to cost.

3.2 TYPES OF SENSITIVITY ANALYSIS

The MITRE NRLA program will perform a sensitivity analysis based on any specified set of changes to LRU/SRU unit cost and MTBF values. The NRLA program user need only list all the LRU/SRU index numbers for which sensitivity analysis is desired, together with the desired percentage changes in MTBF and unit cost values for each item. The program will then recalculate the repair level cost factors associated with every affected LRU failure mode, rerun the RLA algorithm, and print out all resulting changes in repair level decisions. The sensitivity analysis input file format requirements are given in section 4.3.

This sensitivity analysis approach allows the user complete flexibility in specifying both the particular LRUs and SRUs to be modified and also the degree to which individual MTBF and unit costs are to be changed. The user may also concatenate several lists of sensitivity analyses in one data file in order to perform several different sensitivity analysis cases in one computer run. Thus an "extremes only" sensitivity analysis may be performed by creating one list of changes representing the lower extreme values for MTBF and unit cost, followed by a second list representing the upper extreme values for MTBF and unit cost.

The NRLA program will also allow "wholesale changes" to be made to all LRUs and SRUs by the use of an "ALL" command in the place of a particular index number. Thus one instruction in the sensitivity analysis data file is sufficient to specify a computer run where all LRUs and SRUs have the same percentages changes made to MTBF and unit cost values, respectively. For example, this instruction could specify that all MTBF values be changed by 150 percent and all unit costs values be changed by 80 percent.

One capability that the NRLA program does <u>not</u> provide which is found in the AFLC NRLA model is for automatic "swept sensitivity." In this calculation the model investigates a specified range of, say, MTBF values, one LRU or SRU at a time, and notes values where any change in repair level decisions occurred. Swept sensitivity requires an intensive amount of computer processing if the basic NRLA data base is even of moderate size. It is not routinely performed by users of the current AFLC NRLA model. However, one can still perform swept sensitivity using the MITRE NRLA program, but one must create an explicit list in the data file of the MTBF and unit cost values to be swept over.

Details regarding the format of the sensitivity analysis data file are found in section 6.1.

3.3 SENSITIVITY ANALYSIS APPROACH OF MITRE NRLA PROGRAM

The MITRE NRLA program implements sensitivity analysis in a different manner than done in the AFLC NRLA model. Since the AFLC software incorporates a component LCC model, it can recalculate all the required repair level cost factors that would be impacted by a change in, say, the unit cost of a particular LRU.

However, the MITRE NRLA program accepts these basic repair level cost factors (i.e., the costs assigned to the respective repair level arcs of the network) as direct input factors. As a result, it does not have an independent capability to recalculate these costs directly from LCC model equations. (Note that if the user does have an LCC model for use in generating inputs to the NRLA program, then sensitivity analysis could be performed by rerunning this LCC model with the desired modifications to selected parameters, and then feeding the NRLA program with new cost factors to determine the impact on repair level decisions.)

Nevertheless, the MITRE program can still perform sensitivity analysis without being provided a new set of input cost factors. To get around this limitation, the MITRE program requires that each input cost factor be broken out into five components. If we use the variable TIC to stand for the total item cost for a particular LRU/SRU repair decision, then the five components are defined as follows:

CCCF = Component of TIC which is constant with respect to (w.r.t.) both the item unit cost and item reliability

CCLF = Component of TIC which is constant w.r.t. unit cost, but varies linearly w.r.t. reliability (i.e., failure rate)

LCCF = Component of TIC which varies linearly w.r.t. unit cost, but is constant w.r.t. failure rate

LCLF = Component of TIC which changes linearly w.r.t. both unit cost and failure rate

LCSF = Component of TIC which varies linearly w.r.t. unit cost, but which also changes in proportion to the square root of the item failure rate.

(Note: The component names above were chosen to be suggestive, e.g., CCLF can be read as "Constant in Cost, Linear in Failure rate".)

In these definitions by the expression "is constant w.r.t. unit cost" we mean that that cost component will not change if the item unit cost changes. Training costs are an example of a cost which is constant w.r.t. item unit cost. By the expression "varies linearly w.r.t. unit cost" we mean that the given cost component will change in direct proportion to any change in item unit cost; i.e., if the item unit cost doubles then that cost component will double. Replacement spares cost is an example of a cost which varies linearly w.r.t. the item unit cost. Similar interpretations apply to costs which are constant or vary linearly w.r.t. item failure rate. There is also one particular cost component, part of Investment Spares, which charges in proportion to the square root of the failure rate, e.g., if the failure rate increases by a factor of 9, this component will change by approximately a factor of 3.

This last property results from the standard Air Force procedure for computing the required number of Investment (or Initial) Spares for a given item at either the base or depot. This procedure consists of computing an average spares pipeline quantity which we'll denote as SPLQ, and then to add a safety stock which will guarantee a specified confidence level against stock depletion. In particular, AFM 67-1, chapter 11, "Stock Control at Bases", stipulates on page 11-13 that

$$\text{Safety Stock} = \sqrt{3*SPLQ}$$

Thus the total number of initial spares, NISP for an item is given by

$$NISP = SPLQ + \sqrt{3*SPLQ}$$

which must be rounded to the nearest integer value. If the item unit cost is denoted as UC then the cost of initial spares, CISP, is calculated as

$$CISP = SPLQ*UC + \sqrt{3*SPLQ}*UC$$

Since the spares pipeline quantity, SPLQ, varies linearly with respect to failure rate, the first cost term in CISP above belongs to the cost component LCLF and the second term belongs to the cost component LCSF. Below, in table 3-1 we show how each of the standard LCC cost elements can be assigned to one of these five cost components of TIC. Thus we will have

$$TIC = CCCF + CCLF + LCCF + LCLF + LCSF$$

Now if we want to perform sensitivity analysis on a particular item with respect to unit cost and/or failure rate, we can calculate the exact impact on TIC. In particular, if the original item unit cost is changed by a fractional amount FRACUC, so that we have

$$\text{new unit cost} = FRACUC*\text{old unit cost}$$

and if the original failure rate is change by a fractional amount FRACFR, so that

$$\text{new failure rate} = FRACFR*\text{old failure rate}$$

then we may compute the resulting updated value of TIC by the formula

Table 3-1. Repair Level Cost Component Breakdown

Logistics Costs	CCCF	CCLF	LCCF	LCLF	LCSF
Technical Data	X				
Maint. Training		X			
Repair Labor		X			
Inventory Management	X				
Repair Material				X	
Packing & Shipping		X			
Initial Spares					
- pipeline quantity				X	
- safety stock					X
Replacement Spares				X	

$$\text{new TIC} = \text{CCCF} + \text{FRACFR}*\text{CCLF} + \text{FRACUC}*\text{LCCF} + \text{FRACUC}*\text{FRACFR}*\text{LCLF} + \text{FRACUC}*\sqrt{\text{FRACFR}}*\text{LCSF}$$

Note that, by using this method, the NRLA program does not even have to know the original values of the item unit cost or failure rate (or MTBF). It only needs to know the fractional changes FRACUC and FRACFR that we wish to investigate.

To maintain consistency with the AFLC NRLA model we actually perform sensitivity analysis on item MTBF, rather than failure rate. However, since the failure rate is proportional to the inverse of MTBF, the following relationship holds:.

$$\text{FRACFR} = \frac{1}{\text{FRACMTBF}}$$

where the desired fractional change in MTBF is denoted as FRACMTBF. The equation above for the "new" TIC value can then still be applied.

For example, suppose for a particular LRU failure mode, we have TIC = \$3,000 for base level repair, where the five cost components CCCF, CCLF, LCCF, LCLF, and LCSF all equal \$600. If we want to investigate the impact of raising the LRU unit cost by 50 percent and increasing the MTBF by 30 percent, then we have FRACUC = 1.5 and FRACMTBF = 1.3. (So that FRACFR = 1/1.3 = .77.) We can then recompute the TIC for base level repair using the above formula to get:

$$\text{new TIC} = 600 + 0.77 * 600 + 1.5 * 600 + 1.5 * 0.77 * 600$$

$$+ 1.5 * \sqrt{0.77} * 600$$

$$= \$3445$$

This new TIC can then be placed as the cost capacity on the network arc for base level repair. After the cost factors for other repair decisions are similarly adjusted, the NRLA algorithm can then determine any possible changes in the repair level decisions.

3.4 REPAIR LEVEL COST COMPONENT BREAKDOWN

In table 3-1 we provide some guidelines as to how typical logistics costs can be assigned to the five cost components required to perform sensitivity analysis. This table is meant to be suggestive rather than definitive, because the actual dependence of these cost categories on item unit costs and failure rates can vary from one Air Force system to another. For example, there could be a fixed part of the training cost, say for initial type 1 training or training equipment, which does not depend on the item failure rates. This part of the training cost might then be more properly assigned to the CCCF component. There could also be additional cost categories that should be included for particular applications. Thus, the NRLA program user must use his best judgement in applying table 3-1 and in assigning additional cost categories to these five components. In this regard, note in table 3-1 that we have not assigned any cost categories to the LCCF component. However, this component could be required in some system applications, so that we have included it in the NRLA program. (In fact, in our Joint STARS application we have used LCCF to cover the cost of "In-flight Spares" which are determined on the basis of one spare per aircraft, rather than by failure rate.)

Another point to keep in mind is that the logistics costs listed in table 3-1 may not all apply to a given repair level cost factor. This correspondence was discussed in section 2. Finally, the detailed format required for data input of each repair level cost factor, and its five cost components, is given in section 4.

SECTION 4

COMPUTER OPERATION OF NRLA PROGRAM

4.1 EXECUTING THE PROGRAM

The NRLA program consists of about 1,100 lines of Pascal language computer code. It was developed on a Digital VAX 11/780 computer under the VMS operating system. The NRLA program requires eight input files, seven of which are required to execute the NRLA baseline computer run. The other input file is used for sensitivity analysis.

The seven baseline input files must be given the names below:

FILE1.DAT	(LRU/SRU depot repair costs)
FILE2.DAT	(LRU/SRU base repair costs)
FILE3.DAT	(LRU scrap costs)
FILE4.DAT	(SRU scrap costs)
FILE5.DAT	(support equipment cross reference file)
FILE6a.DAT	(base support equipment costs)
FILE6b.DAT	(depot support equipment costs)

The sensitivity input file must be called "SENS.DAT".

After the Pascal source code has been compiled and linked, the program is invoked by entering the VAX command "$ RUN NRLA". The program will write its output to a file called "NRLA.OUT".

The NRLA program will accommodate up to 1,000 different LRU/SRU types and 200 support equipment types. Operation on such a full data base would require about 90 minutes of elapsed time for one baseline NRLA run, plus one sensitivity analysis run. However, runs on fewer LRU/SRU and support equipment types will execute much more quickly. For example, with about 100 LRU/SRU items a baseline run, including one sensitivity analysis, took only 2 minutes. In addition, the NRLA program which produced these times was compiled under a "range-checking" option which invokes an error check to ensure that each index used as an array argument falls within the bounds of the associated array dimension allocation. This option increases execution times by about a factor of two. Thus a "NOCHECK" compiler option would cut these execution times in half.

The NRLA program could be converted for operation on an IBM PC or compatible, but this would require some modification of the Pascal code. Execution times would probably increase if the program were converted to a PC.

4.2 INPUT DATA FILES FOR BASELINE NRLA DECISIONS

As discussed in section 1, the NRLA program requires that LCC input data be generated outside the program and provided as direct input by the user. If the user has an appropriate LCC model then the methodology described in section 2.4 may be applied so that seven data files may be generated via four separate LCC model runs to produce the required NRLA model input for all network arcs. These respective files contain costs for the following repair decisions and network arcs:

FILE 1: LRU depot repair/ SRU depot repair (LCC run #1)
FILE 2: LRU base repair/SRU base repair (LCC run #2)
FILE 3: LRU scrap for all SRU failures modes (LCC run #3)
FILE 4: SRU scrap and DEC1, DEC2 costs (LCC run #4)
FILE 5: LRU/SRU/SE cross reference file (LCC run #1)
FILE 6a: Base SE cost file (LCC run #2)
FILE 6b: Depot SE cost file (LCC run #1)

4.2.1 Format of Data Files 1 through 3

Figure 4-1 illustrates the general format of data files 1 through 3. Column number 1, for all records, is always left blank.

The cost data for the repair decision represented by the particular data file and for each LRU failure mode appear in successive "blocks" in the data file, with each block containing either two or three records (i.e., lines) of the file. The first record will merely indicate the type of LRU failure mode represented, using the value 1 to indicate a non-SRU related failure mode, and a value of 2 to specify that the LRU failure mode is associated with a particular SRU. The second record in the block will always contain the LRU costs for the particular failure mode. A third record will be given if, and only if, the failure mode is associated with a particular SRU, in which case the SRU costs will be listed. The exact format and content of these records is as follows.

Record No. 1 contains either a 1 or 2 in column 2, indicating the LRU failure mode type, where

1 = LRU only (i.e., a non-SRU-related failure mode)
2 = LRU failure mode associated with SRU

Record No. 2 contains LRU data and, if required, record No. 3 will contain data for the associated SRU, both in the following format: (An "item" is either the LRU or the SRU)

Description	Column (s)	Format
Blank	1,54,65,76,87,98	
Item Index No.	2 -5	Integer, right justified
Item Name	6-29	Alphanumeric, left justified
Item Part No.	30-41	Alphanumeric, left justified
Decision Control Variable	42-43	Integer, right justified

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54
	1																																																				
		1	0	1	L	R	U	T	Y	P	E	1																	8	6	5	0	3	1								0					1	0	0	0	0	0	
	2																																																				
		1	0	2	L	R	U	T	Y	P	E	2																	2	3	6	7	3	2								0						6	5	0	0	0	
		4	0	1	S	R	U	T	Y	P	E	1																	9	8	5	7	3	1								0						2	4	0	0	0	
	2																																																			
		1	0	2	L	R	U	T	Y	P	E	2																	5	6	8	9	5	L								0						7	5	0	0	0	
		4	1	8	S	R	U	T	Y	P	E	24																	6	9	4	3	A									0						3	2	0	0	0	

RECORD NO. 1

COLUMN (S)	DESCRIPTION	FORMAT	JUSTIFICATION
2	FILE CONTENT INDICATOR 1 = LRU DATA ONLY (A NON-SRU FAILURE MODE) 2 = LRU FAILURE MODE ASSOCATED WITH SRU DATA	I1	RIGHT JUSTIFIED

RECORD NO.s 2 AND 3

COLUMN (S)	DESCRIPTION	FORMAT	JUSTIFICATION
1	LEFT BLANK	A1	
2-5	ITEM INDEX NUMBER	I4	RIGHT JUSTIFIED
6-29	ITEM NAME	24A1	LEFT JUSTIFIED
30-41	ITEM PART NUMBER	12A1	LEFT JUSTIFIED
42-43	CONTROL VARIABLE (0 = DECISION ALLOWED,1 = DECISION EXCLUDED)	I2	RIGHT JUSTIFIED
44-53	TIC = TOTAL COST	INTEGER	FREE FORMAT

Figure 4-1. Format for Input Files 1, 2, and 3

55	56	57	58	59	60	61	62	63	64	65	66	67	68	69	70	71	72	73	74	75
					2	5	0	0	0							6	0	0	0	0
					1	8	0	0	0							2	7	0	0	0
						5	0	0	0							1	0	0	0	0
					2	5	0	0	0							1	5	0	0	0
					1	2	0	0	0								5	0	0	0

76	77	78	79	80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95	96	97
						1	5	0	0	0							3	0	0	0	0
						2	0	0	0	0							2	9	0	0	0
							9	0	0	0							1	1	0	0	0
						1	0	0	0	0							1	5	0	0	0
							9	0	0	0								4	0	0	0

98	99	100	101	102	103	104	105	106	107	108	109
						7	0	0	0	0	
						3	6	0	0	0	
						1	3	0	0	0	
						1	0	0	0	0	
							2	0	0	0	

RECORD NO.s 2 AND 3 (CONT.)

COLUMN (S)	DESCRIPTION	FORMAT	JUSTIFICATION
55-64	CCCF = CONSTANT IN COST AND CONSTANT IN RELIABILITY	INTEGER	FREE FORMAT
66-75	CCLF =CONSTANT IN COST AND LINEAR IN RELIABILITY	INTEGER	FREE FORMAT
77-86	LCCF = VARIES LINEARLY IN COST AND CONSTANT IN RELIABILITY	INTEGER	FREE FORMAT
88-97	LCLF = VARIES LINEARLY IN COST AND LINEARLY IN RELIABILITY	INTEGER	FREE FORMAT
99-108	LCSF= VARIES LINEARLY IN COST AND AS THE SQUARE ROOT IN RELIABILITY	INTEGER	FREE FORMAT

Figure 4-1. Format for Input Files 1, 2, and 3 (Continued)

TIC	44-53	Integer, free format
CCCF	55-64	Integer, free format
CCLF	66-75	Integer, free format
LCCF	77-86	Integer, free format
LCLF	88-97	Integer, free format
LCSF	99-108	Integer, free format

where the Decision Control Variable value indicates:

0 = Decision Allowed
1 = Decision Excluded

Thus, for example, if base repair is not possible for a particular LRU failure mode or SRU, then its Decision Control Variable should be set equal to one in the base repair file (Data File 2).

Also note that all cost inputs to these files require an integer format so that decimal points may not be used. In addition, because of the integer format, the largest cost value accepted is $2 billion. Of course, if the actual costs are larger than this figure, they can always be rescaled, e.g., by giving the cost inputs in thousands of dollars.

TIC = Total Item Cost
CCCF = Component of TIC which is constant with respect to both cost and reliability
CCLF = Component of TIC which is constant in cost and varies linearly with respect to reliability (i.e., failure rate)
LCCF = Component of TIC which varies linearly with respect to unit cost and is constant with respect to failure rate
LCLF = Component of TIC which is linear with respect to both unit cost and failure rate
LCSF = Component of TIC which varies linearly with respect to unit cost and in proportion to the square root of the item failure rate.

As a result of this breakout, we have

$$TIC = CCCF + CCLF + LCCF + LCLF + LCSF$$

Only the total item cost TIC is used in the baseline NRLA computer run. The five cost components of TIC are used for any specified sensitivity analysis runs as described in section 3. These five cost components must appear in data files 1 through 4 (with the exception of SRU costs in file 3 -- see below). All subsequent records repeat this block structure until the LRU failure mode combinations are exhausted.

It is important to note that even though file 3 has the same format as files 1 and 2, it is not assumed to contain any SRU cost information since file 3 should be generated by an LCC model run using an LRU "scrap" decision. Thus, SRU costs would be zeroed out by any LCC model as a result of the LRU scrap decision. As a result, the NRLA program merely skips over each SRU line in file 3. For data file readability and consistency checking, it would be good practice to still list the SRU index number and name in file 3 in the prescribed format. However, no SRU costs, even zeroes, need be entered.

The respective records in file 1, 2, and 3 should represent identical LRU failure modes, and associated SRUs, in exactly the same sequence. This convention implies, in particular, that each of these three files will have the same number of records. Thus, no blank lines should appear at the end of any of these files.

4.2.2 Format of Data File 4

Figure 4-2 illustrates the format for data file 4 which contains the cost information for SRU scrap decisions and the costs associated with DEC1 and DEC2, as described in section 2.4. As shown, file 4 contains costs for SRUs only, in blocks of three records for each different SRU type. Record number 1 in each block contains the basic identifying information for an SRU and the scrap arc costs. The format for this record is the same as for record 2 of files 1 through 3. Records 2 and 3, respectively, contain information for the DEC1 and DEC2 costs associated with the SRU and are similar to record number 1 except that columns 1 through 43 are left blank.

The SRUs that are listed in data file 4 should be exactly the same SRUs that appeared in data file 1, 2, and 3, and in exactly the same sequence. Note that even though LRUs are not explicitly listed in file 4, the SRUs do represent specific LRU failure modes. Also, file 4 does not contain any costs for LRU failure modes which are non-SRU related.

4.2.3 Format of Data File 5

Figure 4-3 illustrates the format for data file 5 which indicates which pieces of support equipment are required for corrective maintenance on each ITEM type I. It cross-references LRU failure modes and SRUs with the index numbers of the specific support equipment required for repair at both the base and depot level.

As with files 1 through 3, file 5 is organized into "blocks" of two or three records for each LRU failure mode. Again, the first record of file 5 contains either a 1 or 2 in column 2 where

1 = LRU only (i.e., a non-SRU related failure mode)
2 = LRU failure mode associated with SRU

Record number 2 contains SE requirements for the specified LRU failure mode under both the base and depot repair options and, if required, record number 3 will contain similar SE requirements for the SRU associated with the given failure mode. Both records contain the following information (all integer format, right justified in each field):

Description	Column (s)
Blank	1
Item Index No.	2-5
BNJA	6-9

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54
		2	0	5	S	R	U	N	A	M	E	5																	5	6	3	4	1	5								0						9	0	0	0	0	
																																																2	8	0	0	0	
																																																1	8	0	0	0	
		2	0	9	S	R	U	N	A	M	E	9																	7	6	4	5	X									0											
																																																4	5	0	0	0	
																																																3	2	0	0	0	
																																																2	2	0	0	0	

RECORD NO. 1

COLUMN (S)	DESCRIPTION	FORMAT	JUSTIFICATION
1	LEFT BLANK	A1	
2-5	SRU INDEX NUMBER	I4	RIGHT JUSTIFIED
6-29	SRU NAME	24A1	LEFT JUSTIFIED
30-41	SRU PART NUMBER	12A1	LEFT JUSTIFIED
42-43	CONTROL VARIABLE (0 = DECISION INCLUDED, 1 = DECISION EXCLUDED)	I2	RIGHT JUSTIFIED
44-53	TIC = TOTAL COST	INTEGER	FREE FORMAT

RECORD NO. 2

COLUMN (S)	DESCRIPTION	FORMAT	JUSTIFICATION
1-43	BLANK		
44-53	DEC1 = TOTAL COST	INTEGER	FREE FORMAT

RECORD NO. 3

COLUMN (S)	DESCRIPTION	FORMAT	JUSTIFICATION
1-43	BLANK		
44-53	DEC2 = TOTAL COST	INTEGER	FREE FORMAT

Figure 4-2. Format for Input File 4

55	56	57	58	59	60	61	62	63	64	65	66	67	68	69	70	71	72	73	74	75	76	76	77	78	79	80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99	100	101	102	103	104	105	106	107	108	109
					1	5	0	0	0							3	0	0	0	0								2	0	0	0	0							2	5	0	0	0								5	0	0	0	
						5	0	0	0							1	0	0	0	0									5	0	0	0								5	0	0	0								3	0	0	0	
						3	0	0	0								7	0	0	0									2	0	0	0								5	0	0	0								2	0	0	0	
					1	0	0	0	0							1	2	0	0	0								1	5	0	0	0								3	0	0	0								3	0	0	0	
						8	0	0	0								9	0	0	0									4	0	0	0								7	0	0	0								4	0	0	0	
						4	0	0	0								3	0	0	0									5	0	0	0								4	0	0	0								6	0	0	0	

RECORD NO. 1 (CONT.)

COLUMN (S)	DESCRIPTION	FORMAT	JUSTIFICATION
55-64	CCCF = CONSTANT IN COST AND CONSTANT IN RELIABILITY	INTEGER	FREE FORMAT
66-75	CCLF = CONSTANT IN COST AND LINEAR IN RELIABILITY	INTEGER	FREE FORMAT
77-86	LCCF = VARIES LINEARLY IN COST AND CONSTANT IN RELIABILITY	INTEGER	FREE FORMAT
88-97	LCLF = VARIES LINEARLY IN COST AND LINEARLY IN RELIABILITY	INTEGER	FREE FORMAT
99-108	LCSF = VARIES LINEARLY IN COST AND AS THE SQUARE ROOT IN RELIABILITY	INTEGER	FREE FORMAT

RECORD NO. 2 (CONT.)

COLUMN (S)	DESCRIPTION	FORMAT	JUSTIFICATION
55-64	D1CCCF = CONSTANT IN COST AND CONSTANT IN RELIABILITY	INTEGER	FREE FORMAT
66-75	D1CCLF = CONSTANT IN COST AND LINEAR IN RELIABILITY	INTEGER	FREE FORMAT
77-86	D1LCCF = VARIES LINEARLY IN COST AND CONSTANT IN RELIABILITY	INTEGER	FREE FORMAT
88-97	D1LCLF = VARIES LINEARLY IN COST AND LINEARLY IN RELIABILITY	INTEGER	FREE FORMAT
99-108	D1LCSF = VARIES LINEARLY IN COST AND AS THE SQUARE ROOT IN RELIABILITY	INTEGER	FREE FORMAT

RECORD NO. 3 (CONT.)

COLUMN (S)	DESCRIPTION	FORMAT	JUSTIFICATION
55-64	D2CCCF = CONSTANT IN COST AND CONSTANT IN RELIABILITY	INTEGER	FREE FORMAT
66-75	D2CCLF = CONSTANT IN COST AND LINEAR IN RELIABILITY	INTEGER	FREE FORMAT
77-86	D2LCCF = VARIES LINEARLY IN COST AND CONSTANT IN RELIABILITY	INTEGER	FREE FORMAT
88-97	D2LCLF = VARIES LINEARLY IN COST AND LINEARLY IN RELIABILITY	INTEGER	FREE FORMAT
99-108	D2LCSF = VARIES LINEARLY IN COST AND AS THE SQUARE ROOT IN RELIABILITY	INTEGER	FREE FORMAT

Figure 4-2. Format for Input File 4 (Continued)

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39
	2																																					
		1	0	1				5	8	0	1	8	0	4	7	2	8	4	9	9	6	2	2															
		3	0	4				6		8	7	4	0	5		6	6	2	6	3		7	8	3	3	1												

RECORD NO. 1

COLUMN (S)	DESCRIPTION	FORMAT	JUSTIFICATION
2	FILE CONTENT INDICATOR 1 = LRU DATA ONLY (A NON-SRU FAILURE MODE) 2 = LRU FAILURE MODE ASSOCIATED WITH SRU DATA	I1	RIGHT JUSTIFIED

RECORD NO.S 2 AND 3

COLUMN (S)	DESCRIPTION	FORMAT	JUSTIFICATION
1	LEFT BLANK	A1	
2-5	ITEM INDEX NUMBER	I4	RIGHT JUSTIFIED
6-9	BNJA = NO. DIFFERENT SE TYPES REQUIRED FOR BASE REPAIR	I4	RIGHT JUSTIFIED
10-39	SE INDEX NUMBERS (UP TO 10 NUMBERS OF 3 DIGITS EACH)	9I3	RIGHT JUSTIFIED IN EACH 3 COLUMN FIELD

Figure 4-3. LRU/SRU SE Cross-Reference Input File 5

40	41	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63	64	65	66	67	68	69	70	71	72
			3		2	7	1	1	4		9	7																				
			2		5	4		7	7																							

RECORD NOS. 2 AND 3 (CONT.)

COLUMN (S)	DESCRIPTION	FORMAT	JUSTIFICATION
40-43	DNJA = NO. DIFERENT SE TYPES REQUIRED FOR DEPOT REPAIR	I4	RIGHT JUSTIFIED
44-73	SE INDEX NUMBERS (UP TO 10 NUMBERS OF 3 DIGITS)	9I3	RIGHT JUSTIFIED IN EACH 3 COLUMN FIELD

Figure 4-3. LRU/SRU SE Cross-Reference Input File 5 (Continued)

Base SE index numbers (up to ten #'s of three digits each)	10-39
DNJA	40-43
Depot SE index numbers (up to ten #'s of three digits each)	44-73

where

BNJA = Number of different SE types required for base repair of the item

DNJA = Number of different SE types required for depot repair of the item

The index numbers of the actual SE required for repair at the base and depot are listed, respectively, after the quantities BNJA and DNJA on each record. The records in file 5 should represent the same LRU failure modes, and associated SRUs, in exactly the same sequence as they appear in files 1, 2, and 3.

4.2.4 Format of Data File 6a

Data file 6a (in figure 4-4) contains costs for base level SE, where each record contains the following information; as also shown in figure 4-5:

Description	Column (s)	Format
Blank	1	
SE Index No.	2-5	Integer, right justified
SE Name	6-27	Alpha, left justified
Base Level Cost	28-37	Integer, free format

Data file 6a must contain cost information for all SE types which are identified in data file 5 as being required for base repair of some LRU failure mode or SRU. Omission of any such SE type will cause program termination. On the other hand, SE types which are not required for base repair of any item may be left out of file 6a.

4.2.5 Format of Data File 6b

Data file 6b (see figure 4-5) has a similar format to file 6a, but contains costs for depot level SE, where each record contains the following information:

Description	Column (s)	Format
Blank	1	
SE Index No.	2-5	Integer, right justified
SE Name	6-27	Alpha, left justified
Depot Level Cost	28-37	Integer, free format

Data file 6b must contain cost information for all SE types which are identified in data file 5 as being required for depot repair of some LRU failure mode or SRU. Omission of any such SE

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49
		2	0	1	S	E	N	A	M	E	X	X	X	X	1																	7	0	0	0	0												
		2	0	2	S	E	N	A	M	E	X	X	X	X	2																	4	5	0	0	0												

RECORD NO.s 1 - ...

COLUMN (S)	DESCRIPTION	FORMAT	JUSTIFICATION
1	LEFT BLANK	A1	
2-5	SE INDEX NUMBER	I4	RIGHT JUSTIFIED
6-27	SE NAME	22A1	LEFT JUSTIFIED
28-37	DEPOT LEVEL COST	INTEGER	FREE FORMAT

Figure 4-4. Support Equipment Cost at Base Level (Input File 6a)

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49
		2	0	1	S	E	N	A	M	E	X	X	X	X	1																	6	5	0	0	0												
		2	0	2	S	E	N	A	M	E	X	X	X	X	2																	3	1	0	0	0												

RECORD NO.s 1 - ...

COLUMN (S)	DESCRIPTION	FORMAT	JUSTIFICATION
1	LEFT BLANK	A1	
2-5	SE INDEX NUMBER	I4	RIGHT JUSTIFIED
6-27	SE NAME	22A1	LEFT JUSTIFIED
28-37	DEPOT LEVEL COST	INTEGER	FREE FORMAT

Figure 4-5. Support Equipment Cost at Depot Level (Input File 6b)

type will cause program termination. On the other hand, SE types which are not required for depot repair of any item need not be included in file 6b.

4.3 SENSITIVITY ANALYSIS INPUT FILE

The NRLA program provides the capability to perform sensitivity analysis on repair level decisions, as discussed in section 3. The user may specify a set of LRUs and SRUs and indicate any desired changes to their respective mean time between failures (MTBF) and unit costs (UC). The NRLA program will use these percentage MTBF and unit cost changes to recalculate the network arc costs of all associated LRU failure modes and SRUs, using the methodology described in section 3. By again executing the network algorithm, the program will then find the new optimum repair policy and report any changes from the baseline policy. The paragraph below describes how to create the input file to specify sensitivity runs.

The sensitivity file should begin with the keyword 'RUN'. The text on the same line as 'RUN' is printed in the NRLA output file in order to identify the sensitivity run. The user then lists the LRUs and SRUs whose costs are to be changed on separate lines. The syntax is

{LRU or SRU index number (integer)} {MBTF (integer)} {UC (integer)}

These three numbers may appear anywhere within the first 128 characters on the line and must be separated by at least one blank character. The numbers MTBF and UC here are entered as percentages of the baseline values. For example, 50 indicates half the baseline MBTF, or half the UC.

As a shorthand option, the user may indicate that all LRUs and SRUs should be changed by entering

ALL {MBTF} {UC}

The keywork 'ALL' will cause the network arc costs on all LRU failure nodes and SRUs to be changed. After specifying the changes to be made, the user should enter the line

END

The keyword 'END' signifies that all changes have been entered and that the run should now be made.

The sequence {RUN}...{change specifications}...{END} may be repeated as needed to specify any number of sensitivity runs.

Example

The example file below calls for three sensitivity runs. In the first run, the MBTF of item number 101 is decreased to 50 percent of the baseline value, and its UC is increased to 150 percent of the baseline value. In the second run, these same MTBF and UC changes are applied to all items. In the third run, two different items have modified MTBF and UC factors.

```
RUN   This is run number 1
      101 50 150
END
RUN   This is run number 2.  Change all LRUs and SRUs
      ALL 50 150
END
RUN   This is run number 3
      50 80 120
      80 110 90
END
```

4.4 NRLA OUTPUT

The NRLA program places its output in a file called 'NRLA.OUT'. A sample version of this output file is shown in appendix C. The report first lists the depot support equipment that is recommended for purchase. The section begins with the line 'Depot Support Equipment Purchased', followed by a list giving the support equipment index number, the name, and the cost of the required depot support equipment. At the end of the list, the total cost of the depot support equipment is written. The report then lists the depot support equipment types that are not recommended for purchase. This list gives only the index numbers and names.

The report next lists the recommended and unrecommended base support equipment types in a format identical to the depot support equipment.

The report then lists the recommended repair level decisions for all LRU/SRU failure modes. The report gives the LRU or SRU index number and the item name. In contrast to the data input files 1 through 3, each LRU index number appears only once in the NRLA output file. Below each LRU index number and name, the various LRU failure modes are listed with the SRU-related failure modes being identified by the SRU index number and name. A non-SRU-related failure mode is identified at the bottom of the list by the word "NONE". The repair level decisions for both the LRU (in the given failure mode) and the associated SRU (if any) appear on the same line as the failure mode, with only an LRU decision appearing on the "NONE" line. However, if there are no SRU-related failure modes, then there is only one LRU repair level decision to be made, and this is indicated on the same line as the LRU index number and name (i.e., the "NONE" label is not used).

In this output report there are eight columns headed by the words 'LRU DEPOT', 'LRU SCRAP', 'LRU BASE', 'SRU DEPOT', 'SRU SCRAP', 'SRU BASE', 'DEC1', and 'DEC2'. The recommended repair level is indicated by writing the cost of the decision in the appropriate columns. For example, if the recommended decision is to repair the LRU at the base level and to depot repair the SRU, the columns 'LRU BASE', and 'SRU DEPOT' will have cost numbers in them. Since making this particular decision requires DEC1 and DEC2 costs, the columns headed by 'DEC1' and 'DEC2' will also contain numbers. All other column entries on this line will be blank.

Following the list of recommended decisions, total depot repair, scrap, base repair, DEC1, and DEC2 costs are given. The depot and base totals include both LRU and SRU costs. The depot and base support equipment total costs are repeated here also. The final line is the total estimated cost, which is the sum of the total base, scrap, depot, DEC1, DEC2, base support equipment, and depot support equipment costs.

The report then gives the results of the sensitivity analysis runs, if any are done. The run name, which is read from the 'SENS.DAT' input file, and the specified changes to the MTBF and the unit costs are shown. Then, if any changes to a decision occur as a result of the changed costs, the LRU or SRU index number and name are written, followed by the original LRU decision (DEPOT, BASE, or SCRAP), the original SRU decision, and the new LRU and SRU decisions.

4.5 ERROR MESSAGES

The NRLA program makes the assumption that all input files are in the correct format, and no file format error checking is done.

The costs on an LRU scrap arc in a combined LRU/SRU failure mode are computed by subtracting the cost (found in file4.dat) of scrapping the SRU from the cost (found in file3.dat) or scrapping the LRU. (The reason for this subtraction was discussed in section 2.3.) If the adjusted total item cost is found to be negative, a warning message "warning: negative costs on item <LRU index>, SRU <SRU index>" is written to the screen.

Warnings are also printed to the screen if the total item costs become negative after the adjustments done for sensitivity analysis.

REFERENCES

1. Mitchell, Terrance and Lawrence Briskin, "Network Repair Level Analysis Model: User's Guide," Air Force Logistics Command AFALD/XRS, Wright-Patterson AFB, Ohio, June 1982.

*2. Garvey, R., R. A. Moynihan, K. W. Pullen, and P. Bradley, "Joint STARS Life Cycle Cost Model," ESD-TR-84-201, AD B096 351, Electronic Systems Division, AFSC, Hanscom AFB, MA, August 1985.

*3. Moynihan, Richard A. and Karen W. Pullen, "MILSTAR Terminal Segment Life Cycle Cost Model: User's Manual," MTR-8787 Revision 1, The MITRE Corporation, Bedford, MA, September 1984.

4. "Optimum Repair Level Analysis (ORLA)," AFLCM/AFSCM 800-4, June 1971.

5. Ford, L. R., Jr. and D. R. Fulkerson, *Flows in Networks*, Princeton University Press, Princeton, NJ, 1974.

*These documents have not been reviewed by the Directorate for Security Review and are therefore not available for public dissemination.

GLOSSARY

AFLC	Air Force Logistics Command
DRCT	depot repair cycle time
IMF	intermediate maintenance facility
LCC	life cycle cost
LRU	line replaceable unit
NRLA	Network Repair Level Analysis
MTBF	mean time between failure
RLA	repair level analysis
SE	support equipment
SRU	shop replaceable unit
w.r.t.	with respect to

APPENDIX A

NETWORK FORMULATION

A.1 NETWORK CONSTRUCTION

In order to fully understand the operation of NRLA, it is necessary to understand how the 10 types of decision cost factors can be structured as a network. Figure A-1 a, b, and c shows such a network. Three variations are used in order that all information required may be easily visualized. They all represent the same network. The circles or nodes serve to act as markers defining the arc ends enabling each identification. In terms of a pipeline they represent joints. The lines 1 through 7 represent cuts. These are discussed in section A.3. In section 2, we outlined how the various LRU/SRU repair level decision costs may be calculated and how they relate to the specific network arcs. These cost relationships were shown in a tabularized form in table 2-1 which illustrates how various logistic factors may be summed to form 10 decision cost factors representing the 10 decision components. Figure A-1 represents the 10 cost factors as a network.

A.2 RELATIONAL ARCS

The normal-sized arcs in figure A-1 represent potential decisions. The heavy arcs represent "dummy" or relational arcs. They are used to permit flows but they can never restrain flows. If the network is viewed as a pipeline, arcs 1-2, 1-3, 4-6, 5-4, and 5-6 would have very large diameters, each perhaps with a diameter as large as the sum of all other (non-dummy) arcs. In terms of costs, the very large costs associated with large diameters are pseudo costs. The large pseudo costs prevent the dummy arcs from ever limiting or restricting the flow. We always choose to avoid them in making a network cut (i.e., in trying to block the flow). Thus the line 2a could not be a viable cut. The amount of flow in 1-2 and 1-3 is controlled by the flow in S-1. The flow in 1-2 or 1-3 can never limit the flow in S-1. Notice that there are 10 non-dummy arcs. These represent the 10 possible decision factors discussed in table 2-1.

The 10 decision cost factors when structured as a network with the proper set of relational arcs permit the seven decisions in table 2-2 and the seven cuts of figure A-1.

A.3 CUTS

Before explaining the dummy or relational arcs it is first necessary to define a cut. We should consider that when a set of decisions is selected, the nodes of the network have been divided logically into two sets, those associated with the S, or source, and those associated with the T, or terminal (sink). This is a cut. In appendix B we describe a network "labeling" algorithm which the NRLA program uses to determine the max-flow cut. In figure A-1 we note that each of the cuts individually considered divides the network into two sets of nodes each associated with an S or T.

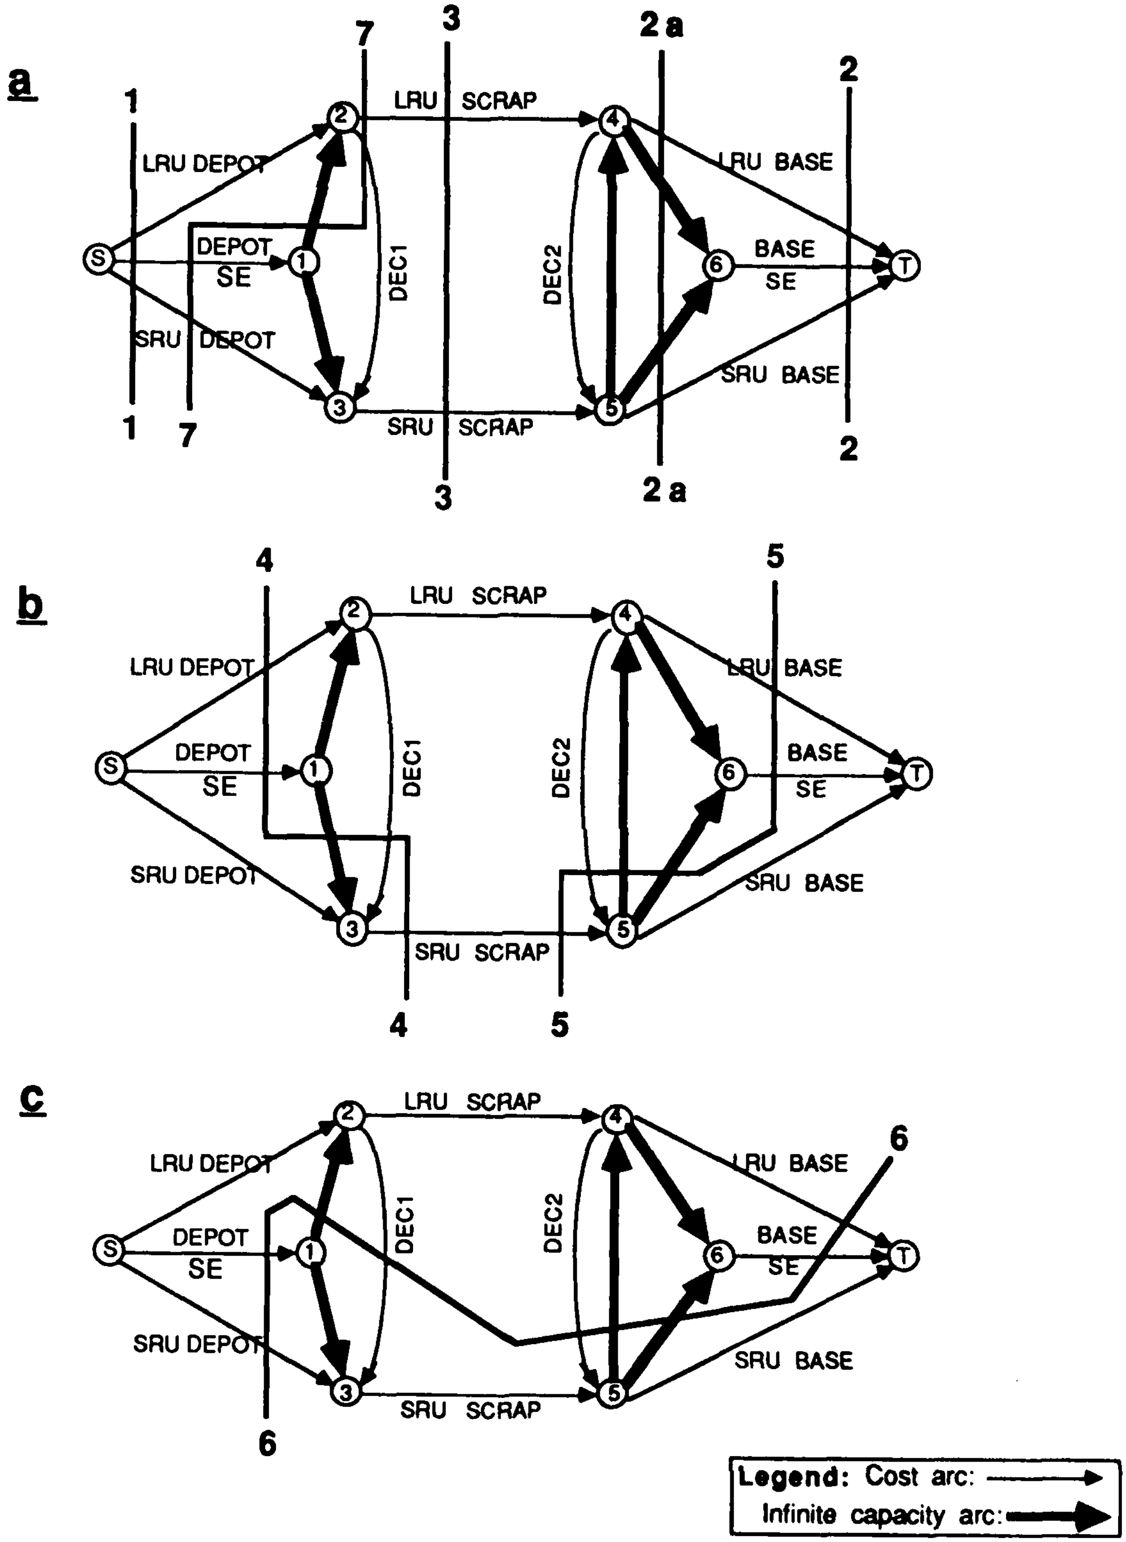

Figure A-1. Sample RLA Network with Cuts

In "max-flow min-cut" problems we wish to find the least cost means of stopping the flow. Once a cut has been selected, all and only those arcs leaving the S set and entering the T set are the ones that "plug" the network. These are the cut set. The reader should consult a book on network theory for greater detail on networks and cuts. (See, for example, *Flows in Networks* [5].)

Referring back to figure A-1 we notice that none of the dummy or relational arcs act as plugs. They all enter the S set or leave the T set. The network was deliberately constructed in this fashion. Their purpose is to prevent certain decisions. As an example: if the line 2a were used as a cut, then the LRU and SRU are both base repaired, but no SE is included. However, the relational arcs 4-6 and 5-6 are so large (costly) that any solution excluding them and other dummies, will be less expensive. Thus, the line 2a would never represent a viable repair decision.

Thus, 2a is a viable cut in that it would separate the network into an S and T set, but since the dummies 4-6 and 5-6 must be counted in the cost of the decision, it is a very expensive decision, and cut 2 will always be a better selection. Thus, the dummies 4-6 and 5-6 force the use of SE at base if either LRU or SRU is base repaired. Arc 5-4 serves to prevent depot repair of the LRU combined with base repair of the SRU, an illogical combination.

A.4 DECISION ARCS

The arcs 2-3 and 4-5 (DEC1 and DEC2) are used in order that certain costs may be included with certain decisions. If the LRU is base repaired and the SRU is scrapped, replacement SRUs must be transported from depot to base and stocked at base. These are DEC2 costs. If the LRU is base repaired and the SRU is depot repaired, SRUs must be stocked at both depot and base and transported both ways. Since DEC2 already carries base stock and one-way transportation costs, DEC1 carries the increased depot stock costs (over that required in an LRU/SRU depot repair situation) and one-way transportation costs. Combined as in cut 6 they represent round trip transportation and stockage at base and depot.

A.5 SCRAP COSTS

Table 2-1 shows how the logistic costs are summed as necessary to obtain the 10 arc costs. This is true for all items except LRU scrap costs. The user should still calculate and input these 10 network arc costs as shown in table 2-1 and as described in section 2. However, the NRLA program will actually use the difference between the input LRU scrap arc cost and the SRU scrap arc cost for its internal LRU scrap arc cost capacity. This difference is used since cut 3, which is the cost of scrapping the LRU, involves arc 2-4, the LRU scrap cost, and arc 3-5, the SRU scrap cost. If summed as described, then the SRU cost would have been included twice, once as part of the LRU scrap cost in arc 2-4 and once as the SRU scrap arc 3-5. By subtracting out the cost of the SRU from the LRU arc 2-4, cut 3, which requires the sum of 2-4 and 3-5, totals to the cost of the LRU when scrapped.

A.6 SOLVING THE SAMPLE NETWORK

Returning to the non-dummy arcs, the costs associated with the 10 respective decision components are used to represent the diameters of the respective arcs. If the cost of stopping the flow in an arc (purchasing a plug) is proportional to its diameter, the question is, what is the minimum cost means of stopping the flow in the network? For which arcs should we buy plugs? It is precisely this problem that the mathematical algorithm called "max-flow min-cut" solves. This total then is the system cost of the decision. As an example in figure A-1, cut 6 selects LRU base repair, base SE, DEC2, DEC1, SRU at depot and depot SE for the failure mode represented. These costs then would be totaled to get the system costs.

If S-2, 2-4, and 4-T represent depot repair, scrap, and base repair for an LRU, respectively, then S-3, 3-5, 5-T represent the same for an SRU, and 2-3, 4-5 represent certain transportation and spares costs. There are seven possible valid decisions. These are shown in table 2-2 and figure A-1.

As an example, cut 4 shows LRU at depot, SE at depot, and SRU scrap being selected. Arc 2-3 (DEC1) is not cut, since it <u>enters</u> the source group of nodes. The total cost for the life cycle is thus determined by summing the three elements: (a) LRU at depot, (b) depot SE, and (c) scrapping the SRU. Table 2-2 shows the same decision in a tabularized form.

By careful review of tables 2-1 and 2-2 and figure A-1, the user will see how the various logistic factors discussed in section 2 may be summed to 10 decision components. These 10 decision components may then be structured as a network if the dummy arcs are used. The network permits only seven feasible decisions. The selection of a particular decision depends upon the costs associated with the arcs of the network. The total cost is the sum of the cost of the cut arcs.

Cut 7 represents an anomalous situation. The LRU is scrapped, but the SRU is depot repaired. It is not possible to structure the network to exclude the possibility of cut 7. If this result occurs, the user should rerun the network twice for each occurrence using a forcing procedure. The runs to be made are (1) excluding LRU scrap and (2) forcing the LRU and SRU to be scrapped. The minimum of these costs represents the optimal solution. (See sections 4.2.1 and 4.2.2 on the use of the "decision control variable" to exclude a given repair level option -- hence to force other options.)

A.7 FULL-SIZE NETWORKS

Where more than one LRU or SRU is involved, additional arcs in parallel to those shown are required. Figure A-2 shows a network with two LRUs, two SRUs, and two types of SE at both depot and at base. Note that figure A-2 specifies that base SE-1 is required by LRU-A, SRU-A, and LRU-B for base level repair. However, SRU-B does not require base SE-1, instead it requires base SE-2.

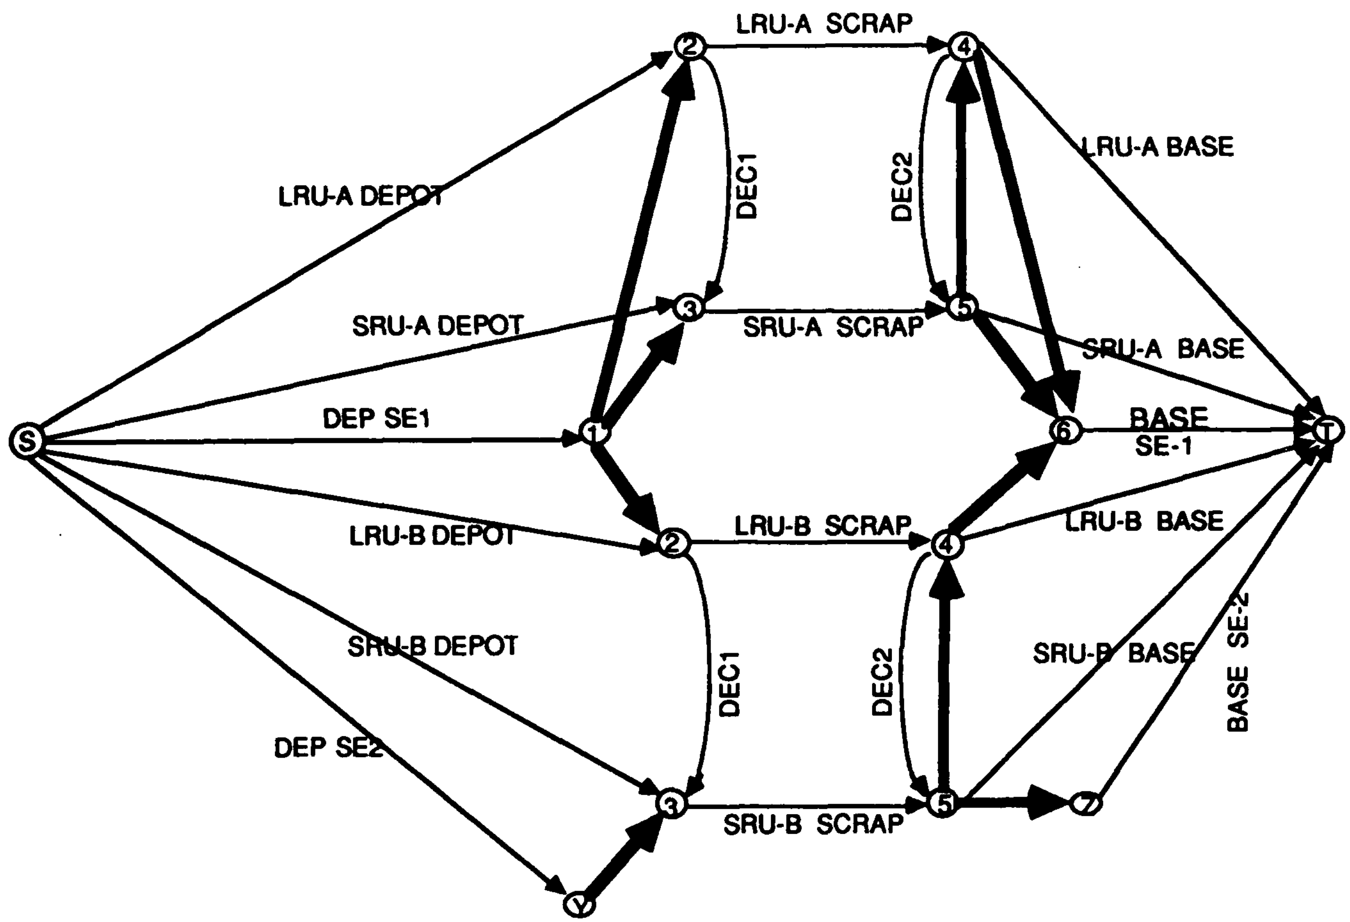

Figure A-2. Sample Multiple LRU/SRU/SE Network

APPENDIX B

NRLA PROGRAM LOGIC AND STRUCTURE

B.1 CONSTRUCTION OF A NETWORK IN NRLA PROGRAM

As discussed in previous sections and in appendix A, a network consists of a set of nodes, and a set of directed arcs connecting the nodes. Each arc of the network that we will construct corresponds a possible decision. Associated with each arc are two numbers. One is the capacity of the arc, which is the cost of the corresponding decision. The other is the flow on the arc, which represents how much we are willing to pay in order to make the decision.

In this section we describe how the NRLA program constructs the network. We begin by making two special nodes, called the source node and the terminal or sink node. The source node will have the property of having no arcs coming into it, and the sink node will have no arcs going out of it.

For each piece of support equipment that might be required at the depot, create a node. Connect an arc from the source node to the created node. Make the capacity of the arc equal to the cost of the depot support equipment. For each piece of support equipment required at the base, create a node. Connect an arc from the new node to the sink node. Make the capacity of the arc equal to the cost of the base support equipment.

For each LRU/SRU failure mode pair, we form the network component shown in figures 2-1 and A-1, as follows: Create four new nodes, called node2, node3, node4, and node5. Connect an arc from the source node to node2, with capacity equal to the cost of depot repair of the LRU (in the given failure mode). Connect an arc from the source node to node3, with capacity equal to the depot repair cost of the SRU. Connect an arc from node2 to node4, with capacity equal to the scrap cost of the LRU minus the scrap cost of the SRU. Connect an arc from node3 to node5, with capacity equal to the scrap cost of the SRU. Connect an arc from node4 to the sink node, with cost equal to the base repair cost of the LRU. Connect an arc from node 5 to the sink node, with cost equal to the base repair cost of the SRU. Connect arcs from node2 to node3, and from node4 to node5, with capacities equal to DEC1 and DEC2, respectively.

For each piece of support equipment required by the LRU (SRU) at the depot, connect an arc from the corresponding support equipment node to node2 (node3), and make the capacity of the arc infinite. Creating these arcs models the constraint that the appropriate support equipment be bought when needed. For each piece of support equipment required by the LRU (SRU) at the base, connect an infinite capacity arc from node4 (node5) to the support equipment node. Finally, connect an infinite capacity arc from node5 to node4. This arc serves to forbid the decisions of scrapping the LRU and base repairing the SRU, and of depot repairing the LRU and base repairing the SRU, neither of which is one of the six feasible decisions.

For each LRU failure mode not associated with any SRU, create two new nodes (2 and 4). Connect arcs from the source node to node2, from node2 to node4, and from node4 to the sink node, with capacities equal to the depot repair, scrap, and base repair costs of the LRU. Make infinite capacity arcs connecting the nodes of depot support equipment needed to node2, and connect infinite capacity arcs from node4 to the required base support equipment nodes.

B.2 THE LABELING METHOD

The network problem as described above is solved by an algorithm called the labeling method. A "label" is a true/false value attached to a node. We call an arc <u>saturated</u> if the flow on the arc equals its capacity; if the flow is less than the capacity we call the arc <u>unsaturated</u>. We call an arc coming out of a node a <u>forward</u> arc, and an arc going into a node a <u>reverse</u> arc. If an arc has a nonzero flow on it, we say that the arc is <u>active</u>. The algorithm is described below.

Initial Condition

All flows in the network are set to zero. All nodes in the network are unlabeled, except the source node, which is labeled.

Labeling rules

If a labeled node has an unsaturated forward arc leading to an unlabeled node, the unlabeled node may now be labeled. If a labeled node has an active reverse arc leading from an unlabeled node, that unlabeled node may also be labeled.

Optimality Condition

If the sink node is unlabeled and no more nodes can be labeled, the algorithm stops. The network now has the maximum possible flow. The corresponding network "cut" passes through all arcs which go from a labeled node to an unlabeled node (but <u>not</u> vice versa). These arcs form the cut set as defined in appendix A.

Breakthrough

If the sink node becomes labeled, we say that <u>breakthrough</u> has occurred. When breakthrough occurs, find the path of arcs connecting the source to the sink node through labeled nodes. This path is called a <u>flow augmenting path</u>. If an arc is a forward arc, the flow on that arc may be increased, up to an amount limited by its capacity. If an arc is a reverse arc, the flow on that arc may be decreased, up to an amount limited by the current flow on the arc. Let L be the minimum amount that flow may be changed along the flow augmenting path. For each forward arc on the path, increase the flow by L. For each reverse arc, decrease the flow by L. The total flow on the network has now been increased by L. Now, make all nodes except the source node unlabeled and continue the process of labeling, getting a breakthrough, and increasing flow until optimality is reached.

Labeling strategy

The description above does not say how to search for nodes to label. There are two obvious strategies. One method is to label all nodes that can be labeled from the currently labeled nodes. Then, label all nodes that can be labeled from the newly labeled nodes, and so on. This is called a <u>breadth first</u> search.

Another method is to label a single node, and then search for the next node to be labeled from the newly labeled node. If no node can be labeled from the current node being searched, back up to the node before it and resume the search. The search stops if the current node is the source node and no more arcs remain to be searched. This is called a <u>depth first</u> search.

The NRLA program as implemented uses a depth first search.

B.3 SOURCE CODE FILES

The source code for the NRLA program is located in five files, which are compiled into four object files. The file 'TYPES.PAS' contains Pascal data type definitions, and is included during the compilation of the other four files.

The file 'NRLA.PAS' contains the code for the routines that do the network algorithm. It also contains the main program, which calls the initialization, optimization, sensitivity analysis, and output procedures. The declarations of all global variables are located here.

The file 'INIT.PAS' contains the code for reading the input files and initializing all data structures.

The file 'OUTPUT.PAS' contains the code for writing the baseline output.

The file 'SENSITIVE.PAS' contains the code that reads the instructions for doing sensitivity analysis, sets up the data structures appropriately, runs the optimization program, and writes the sensitivity analysis output.

B.4 SOURCE CODE DATA STRUCTURES

A network is a set of points called 'nodes' and a set of directed 'arcs' connecting the nodes. The NRLA program contains two arrays called 'network' and 'arcs' to store the nodes and arcs of the network.

Nodes

A node will have a set of arcs coming out of it (forward arcs) and a set of arcs leading into it (reverse arcs). In order to work, the optimization algorithm needs to know

1. The node under consideration. This is held in the global variable 'currentnode'.

2. The arc under consideration. This is held in the field 'currentarc' of the node. Since forward and reverse arcs are treated differently, the algorithm must know the direction of the arc, and this is held in the current node's 'isforward' field, which will be TRUE for a forward arc, and FALSE for a reverse arc.

3. What node to examine if the current node is on a dead end path. This is held in the 'backup' field of the current node.

When a breakthrough occurs, the flow can be increased by an amount equal to the minimum flow change on the arcs on the breakthrough path. For convenience, the minimum on the path up to the current node is stored in the current node's field 'flowup'.

When a node is examined, all arcs coming out of it or leading into it must be examined to see if the node on the other end can be labeled. The program will examine first the forward arcs, then the reverse arcs. The forward and reverse arcs are stored in two single linked lists. The current node has fields 'firstforward' and 'firstreverse' pointing to the beginnings of the respective lists.

Arcs

The arcs in the network repair level analysis formulation correspond to possible repair decisions, or to needed constraints (infinite arcs). All cost information is stored in the array 'arcs', which is an array of records of the type 'arc'. Arc records have fields for the quantities TIC, CCCF, CCLF, LCCF, LCLF, and LCSF which are read out of input data files 1 through 4. The field 'allowed' is set to FALSE if the decision is excluded. The field 'name' is an array of characters that contains the item index number, the item name, and the item part number. Other fields of the arc are:

headnode, tailnode: the nodes at the beginning and end of the arc

nextforward, nextreverse: links for the lists mentioned in the node description

arcflow, capacity: used in the labeling algorithm. The capacity is set to TIC for the baseline run.

decision: tells what kind of arc this arc is. It may take the values infinite, lrudepot2, lruscrap2, lrubase2, srudepot2, sruscrap2, srubase2, lrudepot1, lruscrap1, lrubase1, dec1, dec2, depotse, or basese. Here, the 1 or 2 indicates that the arc is associated with a failure mode of an LRU only (1), or a paired LRU/SRU failure mode (2).

taken, senstaken: the field 'taken' records the decision taken by the baseline run. Senstaken records the decision taken by a sensitivity run.

B.5 PROGRAM UNITS

The NRLA main program is located in the source code file 'NRLA.PAS'. The main program calls the procedures 'initialize', 'optimize', 'outputdata', and 'sensitivity', which are located in the source code files 'INIT.PAS', 'OPTIMIZE.PAS', 'OUTPUT.PAS', and 'SENSITIVE.PAS', respectively.

The procedure 'initialize' reads the seven baseline input files and sets up the data structures needed to do the max-flow, min-cut algorithm accordingly.

The procedure 'optimize' solves the maximum flow problem by using the labeling algorithm described in section B.2.

The procedure 'outputdata' creates a report on the results of the baseline run.

The procedure 'sensitivity' reads the 'SENS.DAT' input file, sets up data structures needed to do the revised max-flow, min-cut algorithm, calls the 'optimize' procedure, and produces a report on the results, which is appended to the same output file that 'outputdata' uses.

APPENDIX C

SAMPLE NRLA RUN

This appendix contains a sample computer run of the NRLA program. The input cost data used represents a fictitious system consisting of three LRU types; a radar, a computer, and a WIU. The radar contains only one SRU, the transmitter, which is responsible for all the radar failures. Thus the radar has only one failure mode, i.e., that due to the transmitter.

The second LRU, the computer, contains three SRUs: a keyboard, a memory module, and a disk drive. However, these three SRUs do not account for all computer failures. Thus the computer has a fourth, non-SRU-related failure mode. The third LRU, the WIU, has no SRUs, so that it has just a single failure mode.

These three LRUs represent all the basically different combinations of failure modes. Thus in figure C-1, the depot repair cost input file, the radar costs are only shown as paired with the transmitter costs, since the radar has no other failure modes. However, the computer costs are broken down and paired with each of its three SRUs in order to list both the LRU and the SRU costs for these SRU-related failure modes. In addition, the computer costs for its non-SRU-related failure mode are separately listed. Finally, the WIU is listed alone since it has no SRUs.

Figures C-2 and C-3 show these LRU and SRU costs in exactly the same format for other repair decisions. However, figure C-4 shows SRU costs only, as discussed in sections 2.4 and 4.2.2. The three lines shown for each SRU represent, respectively, the basic scrap cost (for Replacement Spares), DEC1, and DEC2 costs. Figure C-5 does not show any item names, but gives support equipment requirements for the same LRU/SRU failure mode combinations shown in figures C-1, C-2, and C-3, which are represented in exactly the same sequence. Figures C-6 and C-7 give base and depot costs, respectively, for five sample pieces of support equipment. C-8 shows a simple sensitivity analysis file, as previously discussed in section 4.3.

The sample NRLA output tables are fairly self-explanatory. For example, figure C-9 shows that no depot support equipment was required to be purchased as a result of the NRLA decisions. Figure C-10, however, shows that all base support equipment was needed.

Figure C-11 shows the repair decisions for the baseline NRLA run. In this output table repair decisions are indicated by having their cost printed under the corresponding column heading. Thus on the WIU line in the figure, a cost is shown in the "LRU SCRAP" column indicating the decision for that LRU. Also note in this output table that, for each LRU, all its associated SRUs all appear in a list under the LRU, in contrast to the paired format of figures C-1 through C-3. For each SRU-related failure mode, the repair decision for both the LRU and the SRU is then given on the same line <u>next to the SRU</u>. Thus, for example, on the line next to the memory module, figure C-11 shows that the LRU, the computer, should be base repaired in this failure mode and the SRU, the memory module, should also be base repaired.

In the output table of NRLA decisions, if an LRU has both SRU-related and non-SRU-related failure modes, then the non-SRU-related failure mode will appear last in the list under the LRU and will be identified by the word "NONE" - meaning no SRU is associated with this failure mode. This convention is illustrated in figure C-11 where "NONE" is listed under the computer to indicate its non-SRU-related failure mode. Note that only an LRU repair decision is given for such a failure mode. Finally, if an LRU contains no SRUs, then the repair decision (for the LRU only) is indicated on the same line as the LRU, as is shown for the WIU in figure C-11.

The total system costs for all repair decisions are given at the end of the output baseline NRLA decisions table, broken out into total costs for depot repair, scrap, base repair, DEC1, and DEC2, along with total support equipment costs at both the base and depot levels. Note in figure C-11 that each of the total costs shown for depot repair, scrap, and base repair include both the LRU and SRU costs for the given repair decision. Thus the total base repair cost shown includes the costs for base repair of both the computer (in all its failure modes) and all of its component SRUs.

A final point worth noting in figure C-11, regards the LRU scrap cost shown for the radar (which appears on the line with the transmitter, its only SRU). We have mentioned several times, e.g., in sections 2.3, A.5, and B.1, that, for the LRU scrap network arc cost, the NRLA program actually uses the difference between the input scrap costs of the LRU and the SRU in the given failure mode. This practice avoids double counting SRU scrap costs. Note from figure C-3 that the input scrap cost of the radar is $741,127, and from figure C-4 the input scrap cost of the transmitter is $461,262. Thus figure C-11 shows an LRU scrap cost of $279,865 which is the difference between these two scrap costs. As a result, all the scrap costs shown in figure C-11 can be added together, for both LRUs and SRUs, to get a total system scrap cost without any danger of double counting.

The last figure, C-12 shows the results of a sensitivity analysis run, in which all baseline MTBF values have been changed by 150 percent (i.e., multiplied by 1.5) and all unit costs have been changed by a factor of 50 percent (i.e., cut in half). The output table shows that the repair decisions for only two LRU failure modes were changed, the computer/memory module failure combination and the computer/disk drive failure combination. Note that for these two failure modes, the computer remained as base repaired, but the SRU in each case changed from base repair to scrap. Note that the sensitivity analysis input data file used to generate this run is printed at the beginning of the output results table for reference. (Compare with figure C-8).

```
2
   1  RADAR              1001      0   3812508     41712    3546540        0    224257         0
   2  TRANSMITTER        2002      0  14920902     76212   14778570        0     66120         0
2
   3  COMPUTER           3003      0  84957136     64550   30764988        0  53655548    472049
   4  KEYBOARD           4004      02000000000   3204848 2000000000        0   6922154     18819
2
   3  COMPUTER           3003      0  97321800     73945   35242524        0  61464576    540751
   5  MEMORY MODULE      5005      0 108620016    237519  105851552        0   2526836      4101
2
   3  COMPUTER           3003      0  14406183     10946    5216820        0   9098372     80045
   6  DISK DRIVE         6006      0 618854592    920418  617552640        0    381499         0
1
   3  COMPUTER           3003      0 421769152    320458  152732576        0 266372608   2343484
1
   7  WIU                7007      0   1825030     83407    1741155        0       468         0
```

Figure C-1. Sample LRU/SRU Depot Repair Cost Data File 1

2										
	1	RADAR	1001	0	6510767	43667	6242843	0	224257	0
	2	TRANSMITTER	2002	0	12559337	70816	12422401	0	66120	0
2										
	3	COMPUTER	3003	0	76182120	562243	22549644	0	52711780	358448
	4	KEYBOARD	4004	0	14138338	122201	7066283	0	6918457	31397
2										
	3	COMPUTER	3003	0	87269656	644072	25831518	0	60383448	410616
	5	MEMORY MODULE	5005	0	5604565	154076	2914000	0	2527406	9083
2										
	3	COMPUTER	3003	0	12918202	95340	3823744	0	8938337	60782
	6	DISK DRIVE	6006	0	1747523	143868	1222156	0	381499	0
1										
	3	COMPUTER	3003	0	378205600	2791252	111947552	0	261687264	1779511
1										
	7	WIU	7007	0	3191605	85882	3105255	0	468	0

Figure C-2. Sample LRU/SRU Base Repair Cost Data File 2

2										
	1	RADAR	1001	0	741127	5627	89655	0	645846	0
	2	TRANSMITTER	2002	0	21146	0	21146	0	0	0
2										
	3	COMPUTER	3003	0	174224240	584551	29285724	0	143947008	406967
	4	KEYBOARD	4004	0	21146	0	21146	0	0	0
2										
	3	COMPUTER	3003	0	199580848	669627	33547966	0	164897056	466197
	5	MEMORY MODULE	5005	0	21146	0	21146	0	0	0
2										
	3	COMPUTER	3003	0	29543212	99122	4965981	0	24409098	69009
	6	DISK DRIVE	6006	0	21146	0	21146	0	0	0
1										
	3	COMPUTER	3003	0	864935040	2902001	145388768	0	714623936	2020384
1										
	7	WIU	7007	0	29052	4346	23421	0	1284	0

Figure C-3. Sample LRU Scrap Cost Data File 3

2	TRANSMITTER	2002	0	461262	4705	8986	0	447571	0
				7088	0	7088	0	3835	0
				7088	0	7088	0	0	0
4	KEYBOARD	4004	0	21963552	187122	4776951	0	16999478	0
				2467861	0	2281350	0	142839	43672
				2485624	0	2281350	0	101039	103235
5	MEMORY MODULE	5005	0	6498987	82259	2359731	0	4056996	0
				147015	0	104458	0	34084	8473
				154417	0	104458	0	24127	25831
6	DISK DRIVE	6006	0	3607008	37527	686879	0	2882602	0
				616806	0	581280	0	24229	11297
				643450	0	581280	0	16744	45426

Figure C-4. Sample SRU Scrap Cost Data File 4

```
2
  1   3 2 2 3 0 0 0 0 0 0    3 2 2 3 0 0 0 0 0 0
  2   2 3 2 0 0 0 0 0 0 0    2 3 2 0 0 0 0 0 0 0
2
  3   5 5 5 2 3 4 0 0 0 0    5 5 5 2 3 4 0 0 0 0
  4   4 2 4 5 3 0 0 0 0 0    4 2 4 5 3 0 0 0 0 0
2
  3   5 5 5 2 3 4 0 0 0 0    5 5 5 2 3 4 0 0 0 0
  5.  4 2 3 5 4 0 0 0 0 0    4 2 3 5 4 0 0 0 0 0
2
  3   5 5 5 2 3 4 0 0 0 0    5 5 5 2 3 4 0 0 0 0
  6   4 1 2 3 5 0 0 0 0 0    4 1 2 3 5 0 0 0 0 0
1
  3   5 5 5 2 3 4 0 0 0 0    5 5 5 2 3 4 0 0 0 0
1
  7   4 1 2 3 5 0 0 0 0 0    4 1 2 3 5 0 0 0 0 0
```

Figure C-5. Sample Support Equipment Cross-Reference File 5

1 PECULIAR 1	1514802
2 INT TEST SET 1	222646
3 INT TEST SET 2	30673
4 MBS	234346
5 ATE	63845

Figure C-6. Sample Base Support Equipment Cost File 6a

1 PECULIAR 1	65861
2 INT TEST SET 1	89058
3 INT TEST SET 2	6816
4 MBS	36053
5 ATE	11823

Figure C-7. Sample Depot Support Equipment Cost File 6b

```
RUN This is a test
  ALL  150 50
END
```

Figure C-8. Sample Sensitivity Analysis Input File

Depot Support Equipment Purchased

item index no	name	cost

Total Depot Support Equipment Cost 0

Depot Support Equipment Not Purchased

item index no	name
1	PECULIAR 1
2	INT TEST SET 1
3	INT TEST SET 2
4	MBS
5	ATE

Figure C-9. Sample Output Depot SE Decisions

Base Support Equipment Purchased

item index no	name	cost
1	PECULIAR 1	1514802
2	INT TEST SET 1	222646
3	INT TEST SET 2	30673
4	MBS	234346
5	ATE	63845

Total Base Support Equipment Cost 2066312

Base Support Equipment Not Purchased

item index no	name

Figure C-10. Sample Output Base SE Decisions

LRU/SRU INDEX	NAME	LRU DEPOT	LRU SCRAP	LRU BASE	SRU DEPOT	SRU SCRAP	SRU BASE	DEC1	DEC2
1	RADAR								
2	TRANSMITTER		279865			461262			
3	COMPUTER								
4	KEYBOARD			76182120			14138338		
5	MEMORY MODULE			87269656			5604565		
6	DISK DRIVE			12918202			1747523		
	NONE			378205600					
7	WIU		29052						

TOTALS	DEPOT	SCRAP	BASE	DEC1	DEC2
	0	770179	576066004	0	0

Depot Support Equipment Total	0
Base Support Equipment Total	2066312
Total Cost	578902495

Figure C-11. Sample Output Baseline LRU/SRU Decisions

```
SENSITIVITY ANALYSIS

RUN This is a test
  ALL  150 50
END

LRU/SRU NAME                         BASELINE DECISIONS       ALTERED  DECISIONS
INDEX                                LRU      SRU             LRU      SRU
  3         COMPUTER
      5     MEMORY MODULE            BASE     BASE            BASE     SCRAP
      6     DISK DRIVE               BASE     BASE            BASE     SCRAP
```

Figure C-12. Sample Output Sensitivity Analysis Results